Logical Methods in Computer Science
Vol. 3 (2:3) 2007, pp. 1–27
www.lmcs-online.org

Submitted Apr. 9, 2006
Published Apr. 26, 2007

AUTOMATA WITH NESTED PEBBLES CAPTURE FIRST-ORDER LOGIC WITH TRANSITIVE CLOSURE

JOOST ENGELFRIET AND HENDRIK JAN HOOGEBOOM

Leiden University, Institute of Advanced Computer Science, P.O.Box 9512, 2300 RA Leiden, The Netherlands
e-mail address: {engelfri,hoogeboom}@liacs.nl

Abstract. String languages recognizable in (deterministic) log-space are characterized either by two-way (deterministic) multi-head automata, or following Immerman, by first-order logic with (deterministic) transitive closure. Here we elaborate this result, and match the number of heads to the arity of the transitive closure. More precisely, first-order logic with k-ary deterministic transitive closure has the same power as deterministic automata walking on their input with k heads, additionally using a finite set of nested pebbles. This result is valid for strings, ordered trees, and in general for families of graphs having a fixed automaton that can be used to traverse the nodes of each of the graphs in the family. Other examples of such families are grids, toruses, and rectangular mazes. For nondeterministic automata, the logic is restricted to positive occurrences of transitive closure.

The special case of $k = 1$ for trees, shows that single-head deterministic tree-walking automata with nested pebbles are characterized by first-order logic with unary deterministic transitive closure. This refines our earlier result that placed these automata between first-order and monadic second-order logic on trees.

1. Introduction

The complexity class $\mathsf{DSPACE}(\log n)$ of string languages accepted in logarithmic space by deterministic Turing machines, has two well-known distinct characterizations. The first one, see e.g., [30] (or Corollary 3.5 of [32]), is in terms of deterministic two-way automata with several heads working on the input tape (and no additional storage). Second, Immerman [33] showed that these languages can be specified using first-order logic with an additional deterministic transitive closure operator – it is one of the main results in the field of descriptive complexity [16, 35]. Similar characterizations of $\mathsf{NSPACE}(\log n)$ hold for their nondeterministic counterparts [33, 34, 61].

So we have a match between two distinct ways of specifying languages (two-way multi-head automata and first-order logic with transitive closure) that both have a natural parameter indicating the relative complexity of the mechanism used. For multi-head automata the parameter is the number of heads used to scan the input; indeed, two heads are more powerful than a single one, as two heads can be used to accept nonregular languages like

2000 ACM Subject Classification: F.1.1, F.4.1, F.4.3.
Key words and phrases: tree-walking automata, pebbles, first-order logic, transitive closure.

LOGICAL METHODS IN COMPUTER SCIENCE
DOI:10.2168/LMCS-3 (2:3) 2007
© J. Engelfriet and H. J. Hoogeboom
Creative Commons

$\{a^n b^n \mid n \in \mathbb{N}\}$, whereas single-head two-way automata only accept regular languages [53, 58]. In fact, it is known that in general $k+1$ heads are better than k even for a single-letter input alphabet [44]. For transitive closure logics, the parameter is the arity of the transitive closure operators used: if $\phi(\overline{x},\overline{y})$ denotes a formula that relates two sequences $\overline{x},\overline{y}$ of k variables each, then $\phi^*(\overline{x},\overline{y})$ denotes the transitive (and reflexive) closure of ϕ – we will call this k-ary transitive closure, and it is said to be deterministic if ϕ determines $\overline{y}$ as a function of $\overline{x}$. Clearly, binary transitive closure is more powerful than unary: $\{a^n b^n \mid n \in \mathbb{N}\}$ can easily be described in first-order logic with binary deterministic transitive closure, but first-order logic with unary transitive closure defines the regular languages [2, 16, 52]. It seems to be open whether $(k+1)$-ary transitive closure is more powerful than k-ary transitive closure, see [29].

In [2], Bargury and Makowsky set out to characterize the formulas that capture the power of automata with k heads, and they found a class of formulas, called k-regular, with this property. Apart from first-order concepts, the formulas only use k-ary transitive closure. They show how k-head automata can be described by k-ary transitive closure (both deterministically and nondeterministically) but for the converse the k-ary regular formulas only work in the nondeterministic case: "the modification of the k-regular formulas needed to take out the nondeterminism will spoil their elegant form, and we do not pursue this further" [2].

Here we set out from the other side. Starting with the full set of deterministic k-ary transitive closure formulas we want to obtain an equivalent notion of deterministic automata. Indeed, we succeed in doing this, i.e., we have an automata-theoretic characterization of first-order logic with deterministic k-ary transitive closure, but we pay a price. The two-way automaton model we obtain has k heads, as expected, but is augmented with the possibility to put an arbitrary finite number of pebbles on its input tape, to mark positions for further use. If these pebbles can be used at will it is folklore [54, 51] that we obtain again $\mathsf{DSPACE}(\log n)$, a family too large for our purpose. Instead we only allow pebbles that are used in a LIFO (or nested) fashion: all pebbles can be 'seen' by the automaton as usual, but only the last one dropped can be picked up [26, 19, 43, 65, 22, 49]. On the other hand our pebbles are more flexible than the usual ones: they can be 'retrieved from a distance', i.e., a pebble can be picked up even when no head is scanning the position of the pebble (cf. the "abstract markers" of [4]).

In the nondeterministic case we have to restrict ourselves to formulas with positive occurrences of the k-ary transitive closure operator, as we do not know whether the class of languages accepted by nondeterministic k-head two-way automata with nested pebbles is closed under complement.

In fact, our equivalence result (Theorem 5.3) is stated and proved for ranked trees in general, of which strings are a special case. Our automaton model (a tree-walking automaton) visits the nodes of an input tree, moving up and down along the edges of the tree. The moves (and state changes) are determined by the state of the automaton and the label of (and pebbles on) the nodes it visits; additionally we assume that the children of each node are consecutively numbered and that the automaton can distinguish this number.

In Section 4 we translate logical formulas into automata, following [19] and additionally using the technique of Sipser [59] to deterministically search a computation space. Section 5 considers the reverse: translating automata into logical formulas. As in [2] we adapt Kleene's construction to obtain regular formulas from automata, thus getting rid of the states of the

automaton (Lemma 5.1), but we need to iterate that construction: once for each nested pebble.

In Section 6 we summarize the results of our paper for single-head automata on trees, the context of our previous papers [19, 21] dealing with tree-walking automata and the quest of obtaining simple tree automaton models that are inherently sequential (unlike the classic tree automata) and still capture the full power of regular tree languages.

Finally, in Section 7 we discuss how to extend our results to more general graph-like structures, such as unranked trees (important for XML [43, 65, 47, 37, 57]), cycles, grids (as in [2]; important for picture recognition [4, 56, 25, 41, 40]), toruses, and, for $k \geq 2$, mazes [5, 12, 31]. To have a meaningful notion of graph-walking automaton we only consider graphs with a natural locality condition: a node cannot have two incident edges with the same label and the same direction. Note that unranked trees satisfy this condition when we view them as graphs with 'first child' and 'next sibling' edges in the usual way. Two-dimensional grids satisfy it by distinguishing between horizontal and vertical edges. For all such graphs, our result holds in one direction: from automata to logical formulas. The other direction holds for all families of such graphs for which there exists a single-head deterministic graph-walking automaton (with nested pebbles) that can traverse each graph of the family, visiting each node at least once. This includes all families mentioned above (except mazes, for which the automaton has two heads), and also, trivially, the family of 'ordered' graphs, i.e., all graphs in which the successor relation of a total order is determined by edges with a specific label. Note that the existence of such an automaton is needed to implement even simple logical formulas such as 'all nodes have label σ'.

The main result of this paper, but only for trees and $k = 1$ (cf. Section 6), was first presented at a workshop in Dresden in March 1999, but unfortunately did not make it into the proceedings [66]. Muscholl, Samuelides, and Segoufin [45] have 'reconstructed' our missing result, independently obtaining the closure under complementation of the tree languages accepted by deterministic tree-walking automata with nested pebbles, taking special care to minimize the number of pebbles needed. The results of this paper were then presented at STACS 2006 [20].

2. Preliminaries

Trees. A *ranked alphabet* is a finite set Σ together with a mapping $\mathrm{rank} : \Sigma \to \mathbb{N}$. Terms over Σ are recursively defined: if $\sigma \in \Sigma$ is of rank n, and $t_1, \ldots, t_n$ are terms, then $\sigma(t_1, \ldots, t_n)$ is a term. In particular σ is a term for each symbol σ of rank 0.

As usual, terms are visualized as *trees*, which are special labelled graphs; $\sigma(t_1, \ldots, t_n)$ as a tree which has a root labelled by σ and outgoing edges labelled by $1, \ldots, n$ leading to the roots of trees for $t_1, \ldots, t_n$. The roots of subtrees $t_1, \ldots, t_n$ are said to have child number $1, \ldots, n$, respectively; by default the child number of the root of the full tree equals 0. The set of all trees (terms) over ranked alphabet Σ is denoted by T_Σ.

Strings over an alphabet Σ can be seen as a special case: they form 'monadic' trees over $\Sigma \cup \{\bot\}$, where $\mathrm{rank}(\sigma) = 1$ for each $\sigma \in \Sigma$, and $\mathrm{rank}(\bot) = 0$.

Tree-walking automata. For $k \geq 1$, a k-head *tree-walking automaton* is a finite-state automaton equipped with k heads that walks on an input tree (over a given ranked alphabet

Σ) by moving its heads along the edges from node to node[1]. At each moment it determines its next step based on its present state, and the label and child number of the nodes visited. Accordingly, it changes state and, for each of its heads, it stays at the node, or moves either up to the parent of the node, or down to a specified child. If the automaton has no next step, we say it *halts*.

The *language* $L(\mathcal{A}) \subseteq T_\Sigma$ *accepted* by the k-head tree-walking automaton $\mathcal{A}$ is the set of all trees over Σ on which $\mathcal{A}$ has a computation starting with all its heads at the root of the tree in the initial state and halting in an accepting state, again with all heads at the root of the tree. The family of languages accepted by k-head deterministic tree-walking automata is denoted by $\mathsf{DW}^k\mathsf{A}$, for nondeterministic automata we write $\mathsf{NW}^k\mathsf{A}$.

Such an automaton is able to make a systematic search of the tree (which can be tuned to be, e.g., a preorder traversal), even using a single head, as follows. When a node is reached for the first time (entering it from above) the automaton continues in the direction of the first child; when a leaf is reached, the automaton goes up again. If a node is reached from below, from a child, it goes down again, to the next child, if that exists; otherwise the automaton continues to the parent of the node. The search ends when the root is entered from its last child. This traversal is often used in constructions in this paper.

In both [48] and [50], as an example, the authors explicitly construct a deterministic 1-head tree-walking automaton that evaluates boolean trees, i.e., terms with binary operators 'and' and 'or' and constants 0 and 1.

Again, strings form a special case. Tree-walking automata on monadic trees are equivalent to the usual two-way automata on strings. A tree-walking automaton is able to recognize the root of a tree as well as its leaves (using child number and rank of the symbols). This corresponds to a two-way automaton moving on a tape, where the input string is written with two endmarkers so the automaton knows the beginning and end of its input.

Logic for Trees. For an overview of the theory of first-order and monadic second-order logic on both finite and infinite strings and trees in relation to formal language theory, see [63].

In this paper our primary interest is in *first-order* logic, describing properties of trees. The logic has node variables $x, y, \ldots$, which for a given tree range over its nodes. There are four types of atomic formulas over Σ: $\mathrm{lab}_\sigma(x)$, for every $\sigma \in \Sigma$, meaning that x has label σ; $\mathrm{edg}_i(x, y)$, for every i at most the rank of a symbol in Σ, meaning that the i-th child of x is y; $x \leq y$, meaning that x is an ancestor of y; and $x = y$, with obvious meaning. The formulas are built from the atomic formulas using the connectives $\neg$, $\wedge$, and $\vee$, as usual; variables can be quantified with $\exists$ and $\forall$.

If t is a tree over Σ, ϕ is a formula over Σ such that its free variables are $x_1, \ldots, x_n$, and $u_1, \ldots, u_n$ are nodes of t, then we write $t \models \phi(u_1, \ldots, u_n)$ if formula ϕ holds for t where the free x_i are valuated as u_i.

For fixed $k \geq 1$, by overlined symbols like $\overline{x}$ we denote k-tuples of objects of the type referred to by x, like logical variables, nodes in a tree, or pebbles used by an automaton. By $x[i]$ we then denote the i-th component of $\overline{x}$.

We consider the additional operator of k-ary *transitive closure*. Let $\phi(\overline{x}, \overline{y})$ be a formula where $\overline{x}, \overline{y}$ are k-tuples of distinct variables occurring free in ϕ. We use $\phi^*(\overline{x}, \overline{y})$ to denote the transitive closure of ϕ with respect to $\overline{x}, \overline{y}$. Informally, $\phi^*(\overline{x}, \overline{y})$ means that we can make a series of jumps from k-tuple $\overline{x}$ to k-tuple $\overline{y}$ such that each pair of consecutive k-tuples

[1]Maybe its heads should be called feet.

$\overline{x}', \overline{y}'$ connected by a jump satisfies $\phi(\overline{x}', \overline{y}')$. More formally, let ϕ have $2k+m$ free variables $\overline{x}, \overline{y}, z_1, \ldots, z_m$. For tree t and nodes $\overline{u}, \overline{v}, w_1, \ldots, w_m$ of t we have $t \models \phi^*(\overline{u}, \overline{v}, w_1, \ldots, w_m)$ if there exists a sequence of k-tuples of nodes $\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$, $n \geq 0$, such that $\overline{u} = \overline{u}_0$, $\overline{v} = \overline{u}_n$, and $t \models \phi(\overline{u}_i, \overline{u}_{i+1}, w_1, \ldots, w_m)$ for each $0 \leq i < n$. In particular, $t \models \phi^*(\overline{u}, \overline{u}, w_1, \ldots, w_m)$ for every k-tuple $\overline{u}$ of nodes of t.

Formally we should specify the k-tuples $\overline{x}, \overline{y}$, or rather $\overline{x}', \overline{y}'$, of free variables with respect to which to take the transitive closure, like for the usual universal and existential quantification, and write $(\mathsf{tc}(\overline{x}', \overline{y}')\phi(\overline{x}', \overline{y}'))(\overline{x}, \overline{y})$ instead of $\phi^*(\overline{x}, \overline{y})$.

A predicate $\phi(\overline{x}, \overline{y})$ with free variables $\overline{x}, \overline{y}$ is *functional* (in $\overline{x}, \overline{y}$) if for every tree t and k-tuple of nodes $\overline{u}$ there is at most one k-tuple $\overline{v}$ such that $t \models \phi(\overline{u}, \overline{v})$. If ϕ has more free variables than $\overline{x}, \overline{y}$, this should hold for each fixed valuation of those variables. The transitive closure $\phi^*(\overline{x}, \overline{y})$ is *deterministic* if ϕ is functional (in the variables with respect to which the transitive closure is taken). Instead of requiring ϕ to be functional we could, equivalently, require it to be of the form $\psi(\overline{x}, \overline{y}) \wedge \forall \overline{z}(\psi(\overline{x}, \overline{z}) \rightarrow \overline{y} = \overline{z})$. This has the advantage of being a decidable property, but it is less convenient in proofs.

The *tree language* defined by a closed formula ϕ over Σ consists of all trees t in T_Σ such that $t \models \phi$. The family of all tree languages that are first-order definable is denoted by FO; if one additionally allows k-ary transitive closure or deterministic transitive closure we have the families $\mathsf{FO{+}TC}^k$ and $\mathsf{FO{+}DTC}^k$, respectively. General transitive closure and deterministic transitive closure (i.e., over unbounded values of k) characterize the complexity classes $\mathsf{NSPACE}(\log n)$ and $\mathsf{DSPACE}(\log n)$ (for strings, or more generally, ordered structures), respectively, see [16, 35].

By LFO we denote the family of tree languages definable in *local* first-order logic, i.e., dropping the atomic formula $x \leq y$. One should note however, that $x \leq y$ is the transitive closure of the functional predicate 'x parent of y', i.e., $\bigvee_i \mathrm{edg}_i(x, y)$. Hence $x \leq y$ is expressible in LFO+DTC^1, and the families $\mathsf{FO{+}DTC}^k$, etc., of tree languages definable in first-order logic with transitive closure, do not change by this restriction.

In Section 6 we study the specific case $k = 1$, i.e., we consider unary transitive closure only. The family $\mathsf{FO{+}TC}^1$ is included in the family of regular tree languages, i.e., the tree languages that are definable in *monadic second-order* logic, which additionally has node set variables $X, Y, \ldots$, ranging over sets of nodes of the tree; it allows quantification over these variables, and has the predicate $x \in X$, with its obvious meaning.

Example 2.1. It is proved in [8] that there is a regular tree language T that cannot be accepted by any single-head nondeterministic tree-walking automaton. Here we illustrate how to construct an FO+DTC^1 formula for that language.

Let $\Sigma = \{a, b, c\}$, where a and b are nullary (labelling, of course, the leaves of the trees over Σ) and c is binary (labelling the internal nodes). The language T consists of all trees over Σ for which the path to each leaf labelled by a contains an even number of 'branching' nodes, i.e., internal nodes for which both the left and right subtree contain an a-labelled leaf.

It is easy to construct a first-order formula expressing that a node is branching. Let $\psi(x, y)$ specify that y is the lowest branching ancestor of x, and let $\phi(x, y) \equiv (\exists z)(\psi(x, z) \wedge \psi(z, y))$ to claim y is the second lowest branching ancestor of x. Observe that ϕ is functional. Now T is specified by the FO+DTC^1 formula $(\forall x)(\mathrm{lab}_a(x) \rightarrow (\exists y)[\ \phi^*(x, y) \wedge \neg(\exists z)\psi(y, z)\ ]\ )$.

In fact T even belongs to FO, as observed in [8], but earlier in [52, Lemma 5.1.8]. □

3. Tree-Walking Automata with Nested Pebbles

A k-head tree-walking automaton with *nested pebbles* is a k-head tree-walking automaton that is additionally equipped with a finite set of pebbles. During the computation it may drop these pebbles (one by one) on nodes visited by its heads, to mark specific positions. It may test the currently visited nodes to see which pebbles are present. Moreover, it may retrieve a pebble from anywhere in the tree, provided the *life times* of the pebbles are nested. This can be formalized by keeping a (bounded) stack in the configuration of the automaton, pushing and popping pebbles when they are dropped and retrieved, respectively. Note that since this stack is bounded by the number of pebbles, it can also be kept in the finite control of the automaton.[2] Pebbles can be reused any number of times (but there is only one copy of each pebble). Accepting computations should start and end with all heads at the root without pebbles on the input tree.

The family of tree languages accepted by deterministic k-head tree-walking automata with nested pebbles is denoted by $\mathsf{DPW}^k\mathsf{A}$, the nondeterministic variant by $\mathsf{NPW}^k\mathsf{A}$.

Some specific properties of these pebbles must be stressed. First, as stated above, pebbles are used in a LIFO manner, as in [26, 19, 43, 65, 22, 49], which means that only the last one dropped can be retrieved, and thus their life times on the tree are nested. Without this restriction again the classes $\mathsf{DSPACE}(\log n)$ and $\mathsf{NSPACE}(\log n)$ would be obtained. Second, this is rather nonstandard, the automaton need not return one of its heads to the position where a pebble was dropped in order to pick it up: at any moment the last pebble dropped can be retrieved. This means that the pebble behaves as a *pointer*: we can store the address of a node when we know it (which is the case when we visit it) and we can later wipe the address from memory without the need to return to the node itself. Such pebbles were called "abstract markers" in [4] (to distinguish them from the usual "physical markers"). Finally, as opposed to [26], during the computation all pebbles dropped remain visible to the automaton (and not only the one or two on top of the stack).

Example 3.1. The regular tree language T from Example 2.1 cannot be accepted by any single-head nondeterministic tree-walking automaton (without pebbles), as proved in [8]. As an example, here we show how to accept T by a (single-head) deterministic tree-walking automaton with two nested pebbles.

Using a preorder traversal of the input tree, the first pebble is placed consecutively on leaves labelled by a. For each such position, starting at the leaf we follow the path upwards to the root counting the number of branching nodes. To test whether an internal node is branching we place the second pebble on the node and test whether its other subtree, i.e., the subtree that does not contain the first pebble, contains an a-labelled leaf (using again a preorder traversal of that subtree, the root of which can be recognized through the second pebble which marks the parent of that root). After testing the node, we pick up the second pebble. At the root, we reject whenever we count an odd number of such nodes on a path; otherwise we return to the position of the first pebble (using another preorder traversal of the tree).

[2] If the automaton uses pebbles $x_1, \ldots, x_n$, then the contents of the stack can be any string over $\{x_1, \ldots, x_n\}$ in which each x_i occurs at most once. It can be assumed w.l.o.g. that the stack always contains $x_1 x_2 \cdots x_i$ for some i, but we will do this only in the proof of Lemma 5.2 (where, in fact, the order is reversed).

In fact, the language can be accepted with just one pebble[3] (which is nested trivially). As explained above, the pebble can be used to detect all branching nodes, which, together with all a-labelled leaves, can be viewed as a binary tree. To check, for that tree, that the path to each leaf is of even length, the automaton performs a preorder traversal and counts its number of steps, modulo 2. At each leaf the count should be 0. □

To fix the model, a k-head tree-walking pebble automaton is specified as a tuple $\mathcal{A} = (Q, \Sigma, X, q_0, A, I)$, where Q is a finite set of states, Σ is a (ranked) input alphabet, X a finite set of pebbles, $q_0 \in Q$ the initial state, $A \subseteq Q$ the set of accepting states, and I the finite set of instructions.

Each instruction is a triple of the form $\langle p, \psi, q\rangle$ or $\langle p, \varphi, q\rangle$ or $\langle p, \sim\varphi, q\rangle$, where $p, q \in Q$ are states, ψ is an operation, and φ a test. There are four types of operations: up_i, $\mathrm{down}_{i,j}$ (moves), $\mathrm{drop}_i(x)$, and $\mathrm{retrieve}(x)$ (pebble operations), and three types of tests: $\mathrm{lab}_{i,\sigma}$, $\mathrm{peb}_i(x)$, and $\mathrm{chno}_{i,j}$, where in each case i indicates a head ($1 \le i \le k$), j is a child number ($1 \le j \le \max\{\mathrm{rank}(\sigma) \mid \sigma \in \Sigma\}$), $\sigma \in \Sigma$ is a node label, and $x \in X$ is a pebble. An instruction $\langle p, \chi, q\rangle$ is called an outgoing instruction of state p.

The automaton $\mathcal{A}$ is deterministic if for any pair $\langle p, \chi_1, q_1\rangle$, $\langle p, \chi_2, q_2\rangle$ of distinct instructions starting in the same state, either $\chi_1 = \sim\chi_2$ or $\chi_2 = \sim\chi_1$.

A configuration of $\mathcal{A}$ on tree t over Σ is a triple $[p, \overline{u}, \alpha]$, where $p \in Q$ is a state, $\overline{u}$ is a k-tuple of nodes of t indicating the positions of the k heads, and $\alpha = (x_1, w_1)\cdots(x_m, w_m)$ the stack of pebbles dropped at their positions ($m \ge 0$, $x_j \in X$, w_j a node of t). The initial configuration equals $[q_0, \overline{\mathrm{root}}, \varepsilon]$, where $\overline{\mathrm{root}}$ consists of k copies of the root of t, and ε is the empty stack.

The semantics of the pebble automaton is defined using the step relation $\vdash_{\mathcal{A},t}$ on configurations for automaton $\mathcal{A}$ on input tree t. We have $[p, \overline{u}, \alpha] \vdash_{\mathcal{A},t} [q, \overline{v}, \beta]$ with $\alpha = (x_1, w_1)\cdots(x_m, w_m)$, if there exists an instruction $\langle p, \chi, q\rangle$ such that

	if	then
$\chi =$	up_i	$v[i]$ is the parent of $u[i], v[h] = u[h]$ for $h \ne i, \alpha = \beta$
	$\mathrm{down}_{i,j}$	$v[i]$ is the j-th child of $u[i], v[h] = u[h]$ for $h \ne i, \alpha = \beta$
	$\mathrm{drop}_i(x)$	$\beta = \alpha(x, u[i]), x \notin \{x_1, \ldots, x_m\}, \overline{u} = \overline{v}$
	$\mathrm{retrieve}(x)$	$m \ge 1, \beta(x_m, w_m) = \alpha, x = x_m, \overline{u} = \overline{v}$
	$\mathrm{lab}_{i,\sigma}$	$u[i]$ has label $\sigma, \overline{u} = \overline{v}, \alpha = \beta$
	$\mathrm{peb}_i(x)$	$(x, u[i])$ occurs in $\alpha, \overline{u} = \overline{v}, \alpha = \beta$
	$\mathrm{chno}_{i,j}$	the child number of $u[i]$ is $j, \overline{u} = \overline{v}, \alpha = \beta$

or in case of the negative tests, $\chi = \sim\mathrm{lab}_{i,\sigma}, \sim\mathrm{peb}_i(x), \sim\mathrm{chno}_{i,j}$, the tests above are negated, whereas head positions and pebble stack remain unchanged ($\overline{u} = \overline{v}, \alpha = \beta$).

A configuration c is halting if there is no c' such that $c \vdash_{\mathcal{A},t} c'$, and it is accepting if it is halting and $c = [p, \overline{\mathrm{root}}, \varepsilon]$ for some $p \in A$. Then the language accepted by $\mathcal{A}$ is defined as $L(\mathcal{A}) = \{\, t \in T_\Sigma \mid [q_0, \overline{\mathrm{root}}, \varepsilon] \vdash^*_{\mathcal{A},t} c \text{ for some accepting configuration } c \,\}$.

Example 3.2. Let Σ be as in Example 2.1. We write, in our formalism, a deterministic single-head tree-walking automaton $\mathcal{A}$ (without pebbles) such that $L(\mathcal{A})$ consists of all trees over Σ that have a-labelled leaves only. The automaton performs a preorder traversal of the input tree. The main states are as follows. In state 1 we move down to the left child until we reach a leaf, in state 2 we are at a left child and move to its right sibling, and in state 3 we move up until we are at a left child or at the root.

[3]As brought to our attention by Christof Löding (and his student Gregor Hink).

As the automaton has only a single head, we omit the head number in the instructions. The initial state is 1, the only accepting state is h. The automaton has the following instructions:

$(1, \mathrm{lab}_c, 1')$, $(1', \mathrm{down}_1, 1)$, $(1, \sim \mathrm{lab}_c, 1'')$, $(1'', \mathrm{lab}_a, 3)$,
$(3, \mathrm{chno}_2, 3')$, $(3', \mathrm{up}, 3)$, $(3, \sim \mathrm{chno}_2, 3'')$,
$(3'', \mathrm{chno}_1, 2)$, $(3'', \sim \mathrm{chno}_1, h)$,
$(2, \mathrm{up}, 2')$, $(2', \mathrm{down}_2, 1)$. □

4. From Logic to Nested Pebbles

We now generalize the inclusion $\mathsf{FO} \subseteq \mathsf{DPW}^1\mathsf{A}$ from Section 5 of [19], on the one side introducing k-ary transitive closure, on the other side allowing k heads. Note also the result is for the 'pointer' variant of pebbles, rather than pebbles that have to be picked up where they were dropped.

Lemma 4.1. *For trees over a ranked alphabet, and $k \geq 1$, $\mathsf{FO{+}DTC}^k \subseteq \mathsf{DPW}^k\mathsf{A}$.*

Proof. The proof is by induction on the structure of the formula. For each first-order formula with deterministic transitive closure we construct a deterministic tree-walking automaton with nested pebbles that always halts (with all its heads at the root). The additional effort we have to take to make it always halting, will pay itself back when we deal with negation, but is also helpful when considering disjunction and existential quantification. Generally speaking, each variable of the formula acts as a pebble for the automaton. In case of k-ary transitive closure we need $3k$ pebbles to test the formula by an automaton. Most features can be simulated using a single head, moving pebbles around the tree, only for transitive closure we need all the k heads.

As intermediate formulas may have free variables we need to extend our notion of recognizing a tree by an automaton: a valuation of the free variables is fixed by putting pebbles on the tree, one for each variable, and the automaton should evaluate the formula according to this valuation.

More formally, let $\phi = \phi(x_1, \ldots, x_n)$ be a formula with free variables $x_1, \ldots, x_n$. The automaton $\mathcal{A}$ for ϕ should check whether $t \models \phi(u_1, \ldots, u_n)$ for nodes $u_1, \ldots, u_n$ in a tree t, as follows. It is started in the initial state with all heads at the root of the tree t, where $u_1, \ldots, u_n$ are marked with pebbles $x_1, \ldots, x_n$. During the computation $\mathcal{A}$ may use additional pebbles (in a nested fashion) and it may test $x_1, \ldots, x_n$, but it is not allowed to retrieve them. The computation should halt again with all heads at the root of t with the original configuration of pebbles. The halting state is accepting if and only if $t \models \phi(u_1, \ldots, u_n)$.

For the atomic formulas it is straightforward to construct (single-head) automata. As an example, for $\phi = x \leq y$ the automaton searches the tree for a marked node representing y. From that position the automaton walks upwards to the root, where it halts, signalling whether x was found on the path from y to the root. For $\mathrm{edg}_i(x, y)$ the automaton searches for x, then determines whether x has an i-th child (the arity of the node can be seen from its label) and moves to that child. There the automaton checks whether pebble y is present.

For the negation $\phi = \neg\phi_1$ of a formula we use the original automaton for ϕ_1, but change its accepting states to the complementary set. This construction works thanks to the fact that the automata we build are always halting.

A similar argument works for the conjunction $\phi = \phi_1 \wedge \phi_2$, or the disjunction $\phi = \phi_1 \vee \phi_2$, of two formulas. We may run the two automata constructed for the two constituents consecutively. Note that the free variables in ϕ_1 and ϕ_2 need not be the same, but the extra pebbles that need to be present for ϕ are ignored by each of the two automata.

For quantification $\phi = (\forall x)\phi_1$ the automaton makes a systematic traversal through the tree, using a single head. When it reaches a node (for the first time) it places a pebble x at that position. Then it returns to the root, and runs the automaton for ϕ_1 as a subroutine; the free variable x of ϕ_1 is marked by the pebble, as requested by the inductive hypothesis. When this test for $\phi_1(x)$ is positive, i.e., the subroutine halts at the root in an accepting state, the automaton returns to the node marked x (by searching for it in the tree), picks up the pebble, and places it on the next node of the traversal. When the automaton has successfully run the test for ϕ_1 for each node it accepts. The formula $(\exists x)\phi_1$ is treated similarly.

The main new element of this proof compared to [19] is the introduction of k-ary transitive closure. Here we need to program a walk from one k-tuple of nodes to another k-tuple with 'jumps' specified by a $2k$-ary formula. We cannot do this in a straightforward way, as we might end 'jumping around' in a cycle without noticing. Such an infinite computation violates the requirement that our automaton should always halt. We use a variant of the technique of Sipser [59] to avoid this trap, and run this walk backwards.

So, let $\phi = \phi_1^*$ be the transitive closure of a functional $2k$-ary predicate ϕ_1. Given a tree with $2k$ nodes marked by pebbles $\overline{x}$ and $\overline{y}$ we have to construct an automaton $\mathcal{A}$ that decides whether we can connect the k-tuples $\overline{x}$ and $\overline{y}$ by a series of intermediate k-tuples such that ϕ_1 holds for each consecutive pair. In what follows we assume that ϕ_1 has $2k$ free variables $\overline{x}$ and $\overline{y}$ (with respect to which the transitive closure is taken), and we disregard the remaining free variables of ϕ_1 (the values of which are fixed by pebbles).

Now consider the set of k-tuples of nodes of the input tree t, spanning the (virtual) computation space of $\mathcal{A}$. We build a directed graph on these k-tuples by connecting vertex[4] $\overline{u}$ to vertex $\overline{v}$ if $t \models \phi_1(\overline{u}, \overline{v})$, i.e., the pair $(\overline{u}, \overline{v})$ in t satisfies $\phi_1(\overline{x}, \overline{y})$. As ϕ_1 is functional, for each vertex there is at most one outgoing arc. Thus, if we fix a vertex $\overline{v}$ (of k nodes in t) and throw away the outgoing arc of $\overline{v}$ (if it exists) the component of vertices connected to $\overline{v}$ in this graph forms a tree $t_k(\overline{v})$, with arcs pointing towards the root $\overline{v}$ rather than towards the leaves. Note this is a directed tree in the graph-theoretical sense; there is no bound on the number of arcs incident to each vertex.

Of course, this tree $t_k(\overline{v})$ with ϕ_1-arcs consists of all vertices $\overline{u}$ that satisfy $t \models \phi(\overline{u}, \overline{v})$, and fixing $\overline{v}$ to be the vertex marked by the pebbles $\overline{y}$ the new automaton $\mathcal{A}$ traverses that tree $t_k(\overline{v})$ and tries to find the vertex marked by pebbles $\overline{x}$.

However, the tree $t_k(\overline{v})$ is not explicitly available, and has to be reconstructed while walking on the input tree t, using the automaton $\mathcal{A}_1$ for ϕ_1 as a subroutine. In particular, we want to implement a traversal on $t_k(\overline{v})$. As the vertices of $t_k(\overline{v})$ consist of k-tuples of nodes of t, we order these k-tuples in a natural way using the lexicographical ordering based on the preorder in t. In this way we impose an ordering on the children of each vertex of $t_k(\overline{v})$, thus allowing the usual preorder traversal of $t_k(\overline{v})$ as described below. To find the successor of a k-tuple $\overline{z}$ in the lexicographical ordering we act like adding one to a k-ary number: change the last coordinate of the tuple $\overline{z}$ into its successor (here the preorder

[4]For clarity we distinguish 'node' in the input tree from 'vertex' in the computation space, i.e., a k-tuple of nodes. Similarly we use 'edge' and 'arc'.

successor in t) if that exists, otherwise reset that coordinate to the first element (here the root of t), and consider the last-but-one coordinate, etc.

We implement a preorder traversal of $t_k(\overline{v})$, which means we compute the preorder successor of each vertex of the tree $t_k(\overline{v})$ whose arcs are defined by ϕ_1, and where the ordering between sibling vertices is based on the lexicographical ordering of nodes in t:

preorder successor of vertex $\overline{u}$ in $t_k(\overline{v})$:
    if it exists, the first child of $\overline{u}$,
    else, on the path of $\overline{u}$ to the root $\overline{v}$,
        the right sibling of the first vertex that has one.

We traverse the tree $t_k(\overline{v})$, with $2k$ pebbles $\overline{x}$ and $\overline{y}$ fixed, with the help of $3k$ additional pebbles $\overline{x}'$, $\overline{y}'$, and $\overline{z}'$. During this traversal, $\mathcal{A}$ keeps track of the current vertex of $t_k(\overline{v})$ with its k heads. Initially the heads move to $\overline{y}$, i.e., to $\overline{v}$. Note that the order of dropping the pebbles $\overline{x}'$ and $\overline{y}'$ differs in the two cases below: in the first case we have to check $\phi_1(\overline{x}', \overline{y}')$ 'backwards', finding $\overline{x}'$ given $\overline{y}'$, while in the second case it is the other way around. This is reflected in the order of dropping $\overline{x}'$ and $\overline{y}'$.

First, we describe how to check whether the current vertex has a first child in $t_k(\overline{v})$, and to go there if it exists. We drop pebbles $\overline{y}'$ to fix the current vertex, and we systematically place pebbles $\overline{x}'$ on each candidate vertex, i.e., each k-tuple of nodes of the tree t (except $\overline{v}$). Thus, lexicographically, in each step the last pebble of $\overline{x}'$ is carried to the next node in t (with respect to the preorder in t), but when that pebble has been at all nodes, it is lifted, the last-but-one pebble is moved to its successor node in t, and the last pebble is replaced on the root, etc. For each k-tuple $\overline{x}'$ we check $\phi_1(\overline{x}', \overline{y}')$ using automaton $\mathcal{A}_1$ as a subroutine. If the formula is true, we have found the first child in $t_k(\overline{v})$ and we move the k heads to the nodes marked by $\overline{x}'$, lift pebbles $\overline{x}'$, and retrieve pebbles $\overline{y}'$ (from a distance). If the formula is not true, we move $\overline{x}'$ to the next candidate vertex as described above (but $\overline{v}$ is disregarded). If none of the candidates $\overline{x}'$ satisfies $\phi_1(\overline{x}', \overline{y}')$, the vertex $\overline{y}'$ obviously has no child in $t_k(\overline{v})$.

Second, we describe how to check for a right sibling in $t_k(\overline{v})$, and go there if it exists, or go up (to the parent of the current vertex) otherwise. The problem here is to keep the pebbles in the right order, adhering to the nesting of the pebbles. First drop pebbles $\overline{x}'$ on the current vertex. Then determine its parent in $t_k(\overline{v})$; this is the unique vertex that satisfies $\phi_1(\overline{x}', \overline{y}')$, where $\overline{y}'$ marks the parent vertex of $\overline{x}'$, thanks to the functionality of ϕ_1. It can be found in a traversal of all k-tuples of nodes of t using pebbles $\overline{y}'$ and subroutine $\mathcal{A}_1$ (as described above for the first child). Leave $\overline{y}'$ on the parent and return to $\overline{x}'$ (by searching for $\overline{x}'$ in the tree t). Using the third set of k pebbles $\overline{z}'$, traverse the k-tuples of nodes of t from $\overline{x}'$ onwards and try to find the next k-tuple that satisfies $\phi_1(\overline{z}', \overline{y}')$ when $\overline{z}'$ is dropped. If it is found, it is the right sibling of $\overline{x}'$. Return there, lift $\overline{z}'$, and retrieve $\overline{y}'$ and $\overline{x}'$. If no such k-tuple is found, the current vertex has no right sibling, and we go up in the tree $t_k(\overline{v})$, i.e., we return to $\overline{y}'$. Here we lift $\overline{y}'$ and retrieve $\overline{x}'$.

In all these considerations special care has to be taken of the root $\overline{v}$. It has no parent in $t_k(\overline{v})$. Fortunately $\overline{v}$ is clearly marked by pebbles $\overline{y}$. □

The number of pebbles needed to compute a formula of FO+DTC^k according to the construction above depends only on the nesting of quantifiers and transitive closures in the formula. For each quantifier we count a single pebble, and $3k$ for transitive closure, and compute the maximum needed over all sequences of nested operators in the formula.

When allowing transitive closure of arbitrary formulas (not requiring them to be functional) it is customary to restrict attention to formulas with only *positive occurrences* of transitive closure, i.e., within the scope of an even number of negations (see, e.g., [16, 35]). Using standard argumentation each such formula is equivalent to one where negation is applied to atomic formulas only.

For such formulas there is a similar, nondeterministic, result as the one above. Atomic formulas and their negations are treated as above, and so are conjunction and universal quantification. For disjunction and existential quantification, the automaton uses nondeterminism in the obvious way. For transitive closure, the Sipser technique we have used in the previous proof is not needed. For a formula $\phi = \phi_1^*$ the automaton $\mathcal{A}$ checks nondeterministically the existence of a path $\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$ from vertex $\overline{x}$ to vertex $\overline{y}$ in the directed graph determined by ϕ_1 (described in the proof of Lemma 4.1). When $\mathcal{A}$ is at vertex $\overline{u}_i$, it proceeds to vertex $\overline{u}_{i+1}$ using $2k$ additional pebbles $\overline{x}'$ and $\overline{y}'$, as follows. It drops $\overline{x}'$ on the current nodes $\overline{u}_i$ and nondeterministically chooses nodes $\overline{u}_{i+1}$, where it drops $\overline{y}'$ and checks that $\phi_1(\overline{x}', \overline{y}')$. Then it returns to $\overline{y}'$, lifts $\overline{y}'$, and retrieves $\overline{x}'$.

We denote the positive restriction of $\mathsf{FO{+}TC}^k$ by $\mathsf{FO{+}posTC}^k$, and similarly for the deterministic case. Thus, for trees over a ranked alphabet, nondeterministic k-head tree-walking automata can compute positive k-ary transitive closure: $\mathsf{FO{+}posTC}^k \subseteq \mathsf{NPW}^k\mathsf{A}$.

5. From Nested Pebbles to Logic

The classical result of Kleene [38] shows how to transform a finite-state automaton into a regular expression, which basically means that we have a way to dispose of the states of the automaton. Bargury and Makowsky [2] observe that this technique can also be used to transform multi-head automata walking on grids into equivalent formulas with transitive closure: transitive closure may very well specify sequences of consecutive positions on the input, but has no direct means to store states. A similar technique is used here. As our model includes pebbles, this imposes an additional problem, which we solve by iterating the construction for each pebble. Unlike [2] we have managed to find a formulation that works well for both the nondeterministic and deterministic case.

Given a (deterministic) computational finite-state device with k heads on the tree, the step relation of which is specified by logical formulas, we show that the computation relation that iterates consecutive steps can be expressed using k-ary (deterministic) transitive closure. Of course, the consecutive positions of the heads along the tree are well taken care of by the closure operator, but here we additionally require that the states of the device should match the sequence of steps.

Let Φ be a $Q \times Q$ matrix of predicates $\phi_{p,q}(\overline{x}, \overline{y})$, $p, q \in Q$ for some finite set Q (of states), where $\overline{x}, \overline{y}$ each are k distinct variables occurring free in all $\phi_{p,q}$. We define the *computation closure* of Φ with respect to $\overline{x}, \overline{y}$ as the matrix $\Phi^\#$ consisting of predicates $\phi^\#_{p,q}(\overline{x}, \overline{y})$ where $t \models \phi^\#_{p,q}(\overline{u}, \overline{v})$ iff there exists a sequence of k-tuples of nodes $\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$ and a sequence of states $p_0, p_1, \ldots, p_n$, $n \geq 1$, such that $\overline{u} = \overline{u}_0$, $\overline{v} = \overline{u}_n$, $p = p_0$, $q = p_n$, where $t \models \phi_{p_i,p_{i+1}}(\overline{u}_i, \overline{u}_{i+1})$ for $0 \leq i < n$. ([5])

[5]Note that, to simplify the description of computation closure we have disregarded the remaining free variables of the $\phi_{p,q}$ and $\phi^\#_{p,q}$. More precisely, if $z_1, \ldots, z_m$ are all the free variables of all $\phi_{p,q}$ (in addition to $\overline{x}, \overline{y}$), then each $\phi^\#_{p,q}$ has free variables $\overline{x}, \overline{y}, z_1, \ldots, z_m$. In the definition of $t \models \phi^\#_{p,q}(\overline{u}, \overline{v}, w_1, \ldots, w_m)$ the $z_1, \ldots, z_m$ have fixed values $w_1, \ldots, w_m$.

Intuitively $t \models \phi^{\#}_{p,q}(\overline{u},\overline{v})$ means that there is a Φ-path of consecutive steps (as specified by matrix Φ) leading from nodes $\overline{u}$ in state p to nodes $\overline{v}$ in state q. Note that only nonempty paths are considered ($n \geq 1$).

We say that Φ is *deterministic* if its predicates are both functional and exclusive, i.e., for any $p,q,q' \in Q$ and $3k$ nodes $\overline{u},\overline{v},\overline{v}'$ of any tree t, if both $t \models \phi_{p,q}(\overline{u},\overline{v})$ and $t \models \phi_{p,q'}(\overline{u},\overline{v}')$ then $q = q'$ and $\overline{v} = \overline{v}'$. Moreover, Φ is said to be *semi-deterministic* if the previous requirement holds for final states q,q' only, where q is *final* if $\phi_{q,r}$ is false for all $r \in Q$ (and similarly for q').

Lemma 5.1.

(1) *If Φ is deterministic, then $\Phi^{\#}$ is semi-deterministic.*
(2) *If Φ is in $\mathrm{FO{+}TC}^k$, then so is $\Phi^{\#}$.*
(3) *If Φ is in $\mathrm{FO{+}DTC}^k$ and deterministic, then $\Phi^{\#}$ is in $\mathrm{FO{+}DTC}^k$.*

Proof. 1. Let $q,q' \in Q$ be final, and assume that for tree t both $\phi^{\#}_{p,q}(\overline{u},\overline{v})$ and $\phi^{\#}_{p,q'}(\overline{u},\overline{v}')$ hold for $p \in Q$ and $3k$ nodes $\overline{u},\overline{v},\overline{v}'$ of t, with Φ-paths of length n and n' as in the definition of computation closure ($n,n' \geq 1$).

Consider the first steps of both paths. We have $\phi_{p,p_1}(\overline{u},\overline{u}_1)$ and $\phi_{p,p'_1}(\overline{u},\overline{u}'_1)$, as well as $\phi^{\#}_{p_1,q}(\overline{u}_1,\overline{v})$ if $n \geq 2$, and $\phi^{\#}_{p'_1,q'}(\overline{u}'_1,\overline{v}')$ if $n' \geq 2$. Due to the determinism of Φ we conclude $p_1 = p'_1$ and $\overline{u}_1 = \overline{u}'_1$.

If $n = 1$, then $\overline{u}_1 = \overline{v}$ and $p_1 = q$. Since $p'_1 = q$ is final, $\phi^{\#}_{p'_1,q'}(\overline{u}'_1,\overline{v}')$ is false. Hence $n' = 1$ and $\overline{v}' = \overline{u}_1 = \overline{v}$ and $q' = p_1 = q$ as required. For $n,n' \geq 2$ we continue inductively with p_1 and $\overline{u}_1$.

2. The proof is a logical interpretation of the method of McNaugton and Yamada [42]. Without loss of generality we assume that $Q = \{1,2,\ldots,m\}$. We show by induction on ℓ how to construct a matrix $\Phi^{(\ell)}$ of formulas $\phi^{(\ell)}_{p,q}$ in $\mathrm{FO{+}TC}^k$ which are defined as $\phi^{\#}_{p,q}$, except that the intermediate states $p_1,\ldots,p_{n-1}$ are chosen from $\{1,\ldots,\ell\}$. In particular, for $\ell = 0$ no intermediate states are allowed, whereas for $\ell = m$ all states are allowed, so we have $\Phi^{(m)} = \Phi^{\#}$.

For $\ell = 0$, the length of the path is one. This means that $\Phi^{(0)} = \Phi$.

Given $\Phi^{(\ell)}$ we obtain $\Phi^{(\ell+1)}$ as follows. Assume $\phi^{(\ell+1)}_{p,q}(\overline{x},\overline{y})$ holds. Either there exists a Φ-path that does not visit state $\ell+1$ (i.e., $p_i \neq \ell+1$ for all $0 < i < n$ to be precise), or this state is visited one or more times during the path. In the former case $\phi^{(\ell)}_{p,q}(\overline{x},\overline{y})$ holds, in the latter case we have a path from state p to state $\ell+1$, perhaps looping several times from $\ell+1$ back to itself, and finally there is a path from state $\ell+1$ to state q. Neither of these paths contains $\ell+1$ as intermediate state, so in this case $\phi^{(\ell+1)}_{p,q}(\overline{x},\overline{y})$ postulates the existence of intermediate nodes $\overline{x}'$ and $\overline{y}'$ such that

$$\phi^{(\ell)}_{p,\ell+1}(\overline{x},\overline{x}') \wedge (\phi^{(\ell)}_{\ell+1,\ell+1})^*(\overline{x}',\overline{y}') \wedge \phi^{(\ell)}_{\ell+1,q}(\overline{y}',\overline{y}).$$

3. In the previous part of the proof transitive closure was applied to predicates $\phi^{(\ell)}_{\ell+1,\ell+1}$. However, determinism of Φ entails functionality of predicates of the form $\phi^{(\ell)}_{r,\ell+1}$, by an argument analogous to the one in 1. above. Note that state $\ell+1$ need not be final, but the paths to state $\ell+1$ cannot be extended because (by definition of $\Phi^{(\ell)}$) state $\ell+1$ cannot be visited intermediately. Hence, each transitive closure is applied to a functional predicate, i.e., it is a deterministic transitive closure. □

Lemma 5.2. *For trees over a ranked alphabet, and $k \geq 1$, $\mathsf{DPW}^k\mathsf{A} \subseteq \mathsf{FO{+}DTC}^k$.*

Proof. Consider a tree-walking pebble automaton with k heads. We assume that (1) accepting states have no outgoing instructions (i.e., if $\langle p, \chi, q\rangle$ is an instruction, then p is not accepting), (2) the initial state is not accepting, and (3) if there is an instruction $(p, \mathrm{drop}_i(x), q)$, then there is no instruction $(q, \mathrm{retrieve}(x), r)$. The latter two requirements are to ensure that accepting computations, and computations between dropping and retrieving a pebble, are nonempty, allowing the use of Lemma 5.1.

Let the automaton use n pebbles, $x_n, \ldots, x_1$, where pebbles are placed on the tree in the order given, i.e., x_n is always placed on the bottom of the pebble stack. We view the automaton as consisting of $n{+}1$ 'levels' $\mathcal{A}_n, \ldots, \mathcal{A}_1, \mathcal{A}_0$ such that $\mathcal{A}_\ell$ is a k-head tree-walking pebble automaton with ℓ pebbles $x_\ell, \ldots, x_1$, available for dropping and retrieving, whereas pebbles $x_n, \ldots, x_{\ell+1}$ have a fixed position on the tree and the automaton $\mathcal{A}_\ell$ may test for their presence. Basically, $\mathcal{A}_\ell$ acts as a tree-walking automaton that drops pebble x_ℓ, then queries pebble automaton $\mathcal{A}_{\ell-1}$ with $\ell - 1$ pebbles where to go in the tree, moves there, and retrieves pebble x_ℓ (from a distance).

We postulate that the number of pebbles dropped is kept in the finite control of the automaton, so we can unambiguously partition the state set as $Q = Q_n \cup \cdots \cup Q_1 \cup Q_0$, where Q_ℓ consists of states where ℓ pebbles are still available. The set Q_n contains both initial and accepting states. Automaton $\mathcal{A}_\ell$ equals the restriction of the automaton to the states in Q_ℓ; we will not specify initial and accepting states for $\mathcal{A}_\ell$, $\ell < n$.

We show how to express the computations of automaton $\mathcal{A}_\ell$, $\ell \geq 0$, on the input tree as FO+DTC^k formulas, provided we know how to express computations of automaton $\mathcal{A}_{\ell-1}$ if $\ell \geq 1$. For $\mathcal{A}_\ell$ a matrix $\Phi^{(\ell)}$ is constructed with predicates $\phi^{(\ell)}_{p,q}$ for $p, q \in Q_\ell$. These predicates represent the single steps of $\mathcal{A}_\ell$, so $t \models \phi^{(\ell)\#}_{p,q}(\overline{u}, \overline{v})$ iff $\mathcal{A}_\ell$ has a nonempty computation from configuration $[p, \overline{u}, \alpha]$ to configuration $[q, \overline{v}, \alpha]$. Note that $\Phi^{(\ell)}$ has additional free variables $x_n, \ldots, x_{\ell+1}$ that will hold the positions of the pebbles already placed on the tree, thus representing the pebble stack α.

We first study the steps while the pebble x_ℓ has not been dropped. For each of its heads, automaton $\mathcal{A}_\ell$ may test the presence of one of the pebbles $x_n, \ldots, x_{\ell+1}$, or the node label or the child number of the current node, or it may move the head up to the parent or down to a specified child. The semantics of these separate instructions, relations between the current and next configurations $[p, \overline{u}, \alpha]$ and $[q, \overline{v}, \alpha]$, are easily expressed in first-order logic. So, we have the following translation table:

instruction:	formula:
$\langle p, \mathrm{up}_i, q\rangle$	$\bigvee_j \mathrm{edg}_j(v[i], u[i]) \ \wedge \bigwedge_{h\neq i} u[h] = v[h]$
$\langle p, \mathrm{down}_{i,j}, q\rangle$	$\mathrm{edg}_j(u[i], v[i]) \ \wedge \bigwedge_{h\neq i} u[h] = v[h]$
$\langle p, \mathrm{lab}_{i,\sigma}, q\rangle$	$\mathrm{lab}_\sigma(u[i]) \ \wedge \bigwedge_h u[h] = v[h]$
$\langle p, \mathrm{peb}_i(x_m), q\rangle$	$u[i] = x_m \ \wedge \bigwedge_h u[h] = v[h]$
$\langle p, \mathrm{chno}_{i,j}, q\rangle$	$(\exists u')\, \mathrm{edg}_j(u', u[i]) \ \wedge \bigwedge_h u[h] = v[h]$

or in case of the negative tests $\sim\mathrm{lab}_{i,\sigma}$, $\sim\mathrm{peb}_i(x)$, and $\sim\mathrm{chno}_{i,j}$, the tests above are negated, whereas head positions remain unchanged, e.g., for $\langle p, \sim\mathrm{lab}_{i,\sigma}, q\rangle$ the formula is $\neg\,\mathrm{lab}_\sigma(u[i]) \wedge \bigwedge_h u[h] = v[h]$.

In general $\phi^{(\ell)}_{p,q}(\overline{u}, \overline{v})$ is a disjunction of such formulas, as we may have parallel instructions in the automaton.

Additionally when $\ell \geq 1$, $\mathcal{A}_\ell$ may drop pebble x_ℓ in state p, simulate $\mathcal{A}_{\ell-1}$, and retrieve pebble x_ℓ returning to state q. Such a 'macro step' from configuration $[p, \overline{u}, \alpha]$ to $[q, \overline{v}, \alpha]$ is only possible when there is a pair of pebble instructions $(p, \mathrm{drop}_i(x_\ell), p')$ and $(q', \mathrm{retrieve}(x_\ell), q)$, such that $\mathcal{A}_{\ell-1}$ has a (nonempty) computation from $[p', \overline{u}, \alpha']$ to $[q', \overline{v}, \alpha']$, with $\alpha' = \alpha(x_\ell, u[i])$. Hence, $\mathcal{A}_\ell$ can take a 'step' from $[p, \overline{u}, \alpha]$ to $[q, \overline{v}, \alpha]$ if the disjunction of $\phi^{(\ell-1)\#}_{p',q'}(\overline{u}, \overline{v})$ over all such q' holds, where the free variable x_ℓ in that formula is replaced by $u[i]$, the current position of the i-th head of the automaton, i.e., the position at which that pebble is dropped. Note that in $\mathcal{A}_{\ell-1}$, q' has no outgoing instructions (and hence q' is a final state of $\Phi^{(\ell-1)\#}$).

Defining the remaining $\phi^{(\ell)}_{p,q}$ to be false, we obtain a step matrix $\Phi^{(\ell)}$, which is deterministic thanks to the determinism of the automaton and the semi-determinism of $\Phi^{(\ell-1)\#}$, cf. Lemma 5.1(1). It is in FO+DTC^k by Lemma 5.1(3). The computational behaviour of the automaton $\mathcal{A}_\ell$ is expressed by $\Phi^{(\ell)\#}$, in general, and more specifically for $\mathcal{A}_n$, by the disjunction of all formulas $\phi^{(n)\#}_{p,q}(\overline{\mathrm{root}}, \overline{\mathrm{root}})$ with p the initial state and q an accepting state. Note that the last formula is correct by assumption (1) in the beginning of this proof. □

Combining the two inclusions in Lemma 4.1 and Lemma 5.2, we immediately get the main result of this paper. Note that it includes the case of strings.

Theorem 5.3. *For trees over a ranked alphabet, and $k \geq 1$,* $\mathsf{DPW}^k\mathsf{A} = \mathsf{FO{+}DTC}^k$. □

As a corollary we may transfer two obvious closure properties of $\mathsf{FO{+}DTC}^k$, closure under complement and union, to deterministic tree-walking automata with nested pebbles, where the result is nontrivial. These properties are a rather direct consequence of the always-halting normal form in the proof of Lemma 4.1, which can be obtained for every deterministic automaton. For $k = 1$, this normal form is further studied with regard to the number of pebbles needed in [45, 9].

Corollary 5.4. *Let $k \geq 1$. For each deterministic k-head tree-walking automaton with nested pebbles we can construct an equivalent one that always halts.* □

When the tree-walking automaton is not deterministic we no longer can assure the determinism of the formulas $\Phi^{(\ell)}$ in the proof of Lemma 5.2. However, by Lemma 5.1(2) they are in FO+TC^k.

Theorem 5.5. *For trees over a ranked alphabet, and $k \geq 1$,* $\mathsf{NPW}^k\mathsf{A} \subseteq \mathsf{FO{+}TC}^k$. □

The constructions in the proof of Lemma 5.2 use negation in one place only: it is used on atomic predicates, to model negative tests of the automaton (to check there is no specific pebble on a node). Note that negation is not used to construct the formulas in the proof of Lemma 5.1. Hence we obtain positive formulas, where negation is only used for atomic predicates, and thus the inclusions $\mathsf{DPW}^k\mathsf{A} \subseteq \mathsf{FO{+}posDTC}^k$ and $\mathsf{NPW}^k\mathsf{A} \subseteq \mathsf{FO{+}posTC}^k$. In the deterministic case negation of a transitive closure can (for finite structures) be easily expressed without the negation [28], thus in a positive way: $\mathsf{FO{+}posDTC}^k = \mathsf{FO{+}DTC}^k$. With that knowledge the first inclusion is not surprising; for the nondeterministic case we additionally find a new, positive, characterization (cf. the end of Section 4).

Corollary 5.6. *For trees over a ranked alphabet, and $k \geq 1$,* $\mathsf{NPW}^k\mathsf{A} = \mathsf{FO{+}posTC}^k$. □

As observed in the Introduction, we do not know whether $\mathsf{NPW}^k\mathsf{A}$ is closed under complement (i.e., whether 'pos' can be dropped from Corollary 5.6). Using the method of [34, 61], it is easy to see that, for trees, $\bigcup_{k\in\mathbb{N}} \mathsf{NPW}^k\mathsf{A}$ *is* closed under complement, which means it is equal to $\bigcup_{k\in\mathbb{N}} \mathsf{FO{+}TC}^k$ (and note that it also equals $\bigcup_{k\in\mathbb{N}} \mathsf{NW}^k\mathsf{A}$).

6. Single Head on Trees

More than thirty years ago, single-head tree-walking automata (with output) were introduced as a device for syntax-directed translation [1] (see [23]). Quite recently they came into fashion again as a model for translation of XML specifications [43, 65, 47, 37, 57, 10].

The control of a single-head tree-walking automaton is at a single node of the input tree. Thus it differs from the more commonly known tree automata. These latter automata work either in a top-down or in a bottom-up fashion and are inherently parallel in the sense that the control is split or fused for every branching of the tree.

The power of the classic tree automaton model is well known. It accepts the regular tree languages (both top-down or bottom-up), although the deterministic top-down variant is less powerful. For tree-walking automata however, the situation was unclear for a long time. They accept regular tree languages only [36, 23], but it was conjectured in [18] (and later in [21, 19, 10]) that tree-walking automata cannot accept all regular tree languages[6]. This was first proved for 'one-visit' automata (for the deterministic case in [6, 50], and for the nondeterministic case in [48]). Recently the conjecture was proved, in a very elegant way, for deterministic tree-walking automata in [7], and for nondeterministic tree-walking automata in [8] (see Examples 2.1 and 3.1).

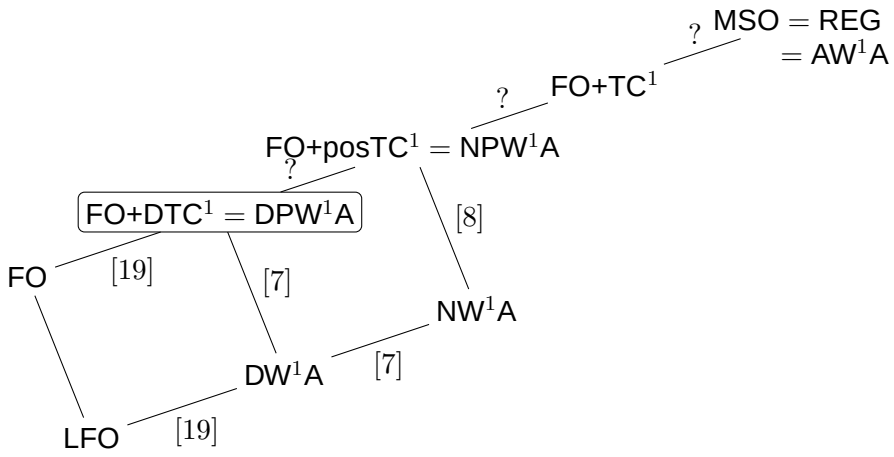
The reason that tree-walking automata cannot fully evaluate trees like bottom-up tree automata is that they easily loose their way. When evaluating a subtree it is in general hard to know when the evaluation has returned to the root of the subtree. In order to facilitate this, in [19] the single-head tree-walking automaton was equipped with pebbles. This was motivated by the ability of pebbles to help finite-state automata find their way out of mazes [5].

In [19] we have shown that all first-order definable tree languages can be accepted by single-head (deterministic) tree-walking automata with nested pebbles, and that tree languages accepted by single-head (nondeterministic) tree-walking automata with nested pebbles are all regular.

As observed before, $\mathsf{DSPACE}(\log n)$ is the class of languages accepted by single-head two-way automata with (nonnested) pebbles [54, 51]. Thus, for $k = 1$ (single-head automata vs. unary transitive closure), our main characterization for tree languages, Theorem 5.3, can be seen as a 'regular' restriction of the result of Immerman characterizing $\mathsf{DSPACE}(\log n)$; on the one hand only (single-head) automata with *nested* pebbles are allowed, while on the other hand we consider only *unary* transitive closure, i.e., transitive closure for $\phi(x, y)$ where x, y are single variables. Note that unary transitive closure can be simulated in monadic second-order logic, which defines the regular tree languages.

We compare the family of tree languages $\mathsf{FO{+}DTC}^1 = \mathsf{DPW}^1\mathsf{A}$ with several next of kin. In the diagram below we have five families of languages $\mathsf{xW}^1\mathsf{A}$ accepted by (single-head) tree-walking automata, which are either deterministic, nondeterministic, or alternating (D,

[6]Although a footnote in [1] claims that the problem was solved by Rabin.



N, or A in x), and may use nested pebbles in case x contains P. Lines without question mark denote proper inclusion, those with question mark just inclusion.

The inclusion LFO $\subseteq$ DW^{1}A was shown in [19], as well as FO $\subseteq$ DPW^{1}A. The regular language $(aa)^*$ cannot be defined in first-order logic, and shows that DW^{1}A $\not\subseteq$ FO. The strictness of DW^{1}A $\subset$ NW^{1}A was shown in [7]; their example additionally shows that FO $\not\subseteq$ DW^{1}A. The result of [8] shows even that FO $\not\subseteq$ NW^{1}A, cf. Example 2.1. Logical characterizations of DW^{1}A and NW^{1}A are given in [48], also using transitive closure (but with an additional predicate indicating the level of a node modulo some constant). All families considered here are contained in the family REG of regular tree languages that can be characterized by monadic second-order logic MSO [15, 62]. The inclusions FO+DTC^1 $\subseteq$ FO+posTC^1 $\subseteq$ FO+TC^1 $\subseteq$ MSO are obvious. In [52] several logics for regular tree languages are studied; it is stated as an open problem whether all regular tree languages can be defined using monadic transitive closure, i.e., whether FO+TC^1 = MSO.

Alternating tree-walking automata are considered in [60]. Alternation combines nondeterminism (requiring a successful continuation from a given state) with its dual (requiring all continuations to be successful). It is not difficult to see that a (nondeterministic) top-down tree automaton can be simulated by an alternating tree-walking automaton, but the reverse inclusion is nontrivial: REG = AW^{1}A.

If, instead of with pebbles, single-head tree-walking automata are equipped with a synchronized pushdown or, equivalently, with 'marbles', then they do recognize all regular tree languages [36, 23, 21], both in the deterministic and nondeterministic case. Synchronization means that the automaton can push or pop one symbol when it moves from a parent to a child or vice versa, respectively.

Questions. Several of the inclusions between the families of trees we have studied are not known to be strict, cf. the figure in this section. These are all left as open problems (but see below). So, for logics, are the inclusions FO+DTC^1 $\subseteq$ FO+posTC^1 $\subseteq$ FO+TC^1 $\subseteq$ MSO strict? For tree-walking automata, are the inclusions DPW^{1}A $\subseteq$ NPW^{1}A $\subseteq$ REG strict, is NW^{1}A $\subseteq$ DPW^{1}A? Considering the use of pebbles, is there a strict hierarchy for tree languages accepted by (deterministic) tree-walking automata in the number of pebbles these automata use?

The pebbles we have used in this paper are considered as pointers, and differ from the pebbles that can only be picked up at the node they are dropped. We conjecture that our type of pebbles is more powerful than the usual one. For nonnested pebbles the two types have the same power, even when the number of pebbles is fixed (as shown in Theorem 2.2 of [4]).

Alternating tree-walking automata with classic (nested) pebbles were considered in [43], where it is shown that these automata accept the regular tree languages. We did not investigate whether alternation together with pointer pebbles again yields the regular languages, i.e., whether $\mathsf{APW}^1\mathsf{A} = \mathsf{REG}$.

Answers. Since the appearance of the report version of this paper, several of the questions above have been answered.

In the recent [9] it is shown that, surprisingly, the two types of nested pebbles have the same power. Moreover, it is shown that the number of pebbles gives rise to a strict hierarchy, both for deterministic and nondeterministic automata, and that the inclusion $\mathsf{NPW}^1\mathsf{A} \subset \mathsf{REG}$ is strict.

In the very recent paper [46] it is shown that $\mathsf{APW}^1\mathsf{A} = \mathsf{REG}$ holds, using an easy variation of the proof technique of [43].

Both tree-walking automata and logical formulas can also be used in a natural way to define binary relations on the nodes of trees, called *trips* in [21]. It is straightforward to show that Theorem 5.3 and Corollary 5.6 also hold for trips. XPath-like formalisms (relevant to XML) that are equivalent with, or closely related to, the trips defined in $\mathrm{FO{+}posTC}^1$ and $\mathrm{FO{+}TC}^1$ have recently been studied in [27, 13].

Single head on strings. For strings the equivalence $\mathsf{MSO} = \mathsf{REG}$ between monadic second-order logic and finite-state automata was obtained independently by Büchi, Elgot, and Trakhtenbrot [11, 17, 64]. As observed, tree-walking automata on monadic trees correspond to two-way finite-state automata on strings. Since the latter are equivalent to ordinary finite-state automata (both deterministic and nondeterministic) [53, 58] most of the hierarchy is known to collapse in the string case: $\mathsf{DW}^1\mathsf{A} = \mathsf{REG}$. Hence for $k = 1$ our main result gives two additional characterizations of the regular string languages $\mathsf{REG} = \mathsf{NPW}^1\mathsf{A}$ and $\mathsf{REG} = \mathsf{FO{+}DTC}^1$. The second equality was shown in [2] (see also Exercise 8.6.3 of [16], and Satz 2.0.1 of [52]).

7. Walking on Graphs

We generalize our results on trees (and strings) to more general families of graphs. There is a basic problem that has to be resolved. We have to propose a model of graphs that is suitable for both graph-walking automata and logic. The structures for which Immerman has obtained his logical characterization of $\mathsf{DSPACE}(\log n)$ are ordered, i.e., they are equipped with a total order $\leq$, which can either be specified directly or, as transitive closure is available in the logical language, by a direct successor relation. In fact, when placing a structure on the tape of a Turing Machine, as has to be done for $\mathsf{DSPACE}(\log n)$, it almost automatically obtains an implicit order on its elements. A natural way to present a structure (assuming only relations of at most arity two) as input to a graph-walking automaton is to represent all pairs in a binary relation as directed edges in a graph, labelled by a symbol representing the relation; unary relations and constants are translated as labels for the nodes. If we fully represent a total order $\leq$ in this way, there will be many edges

with the same label leaving (and entering) each node. Although the edge to the direct successor can be directly specified in first-order logic we see no natural machine model that can choose this edge among the many candidates.

A solution would be to assume the existence of direct successor edges in the graph, i.e., to consider 'ordered' graphs (cf. the end of the Introduction), but we feel this requirement is too restrictive. On trees, tree-walking automata use the preorder to traverse the nodes of the input. Thus, the order is implicit in the tree, and not explicitly given by edges with a special label. Below we generalize this idea to graphs: we do not assume an explicit total order on the nodes, but we postulate the existence of a special graph-walking automaton, the *guide*, that can be used as a 'subroutine' to traverse the nodes of any graph in the family of graphs we are considering. For trees the guide follows the preorder, for (two-dimensional) grids the guide visits the nodes row by row (with generalizations to higher dimensions).

Graphs. We consider (families of) nonempty finite directed graphs, where nodes and edges are labelled using a common alphabet Σ. In order to be able to locally distinguish edges we require that no node has two outgoing edges or two incoming edges with the same label. Hence the graphs we consider are of bounded degree. Parallel edges and loops are allowed. Thus, each edge label represents a (partial) injective function on the set of nodes. We consider (weakly) connected graphs only.

Trees over a ranked alphabet fall under this definition since we label the edge from a parent to its i-th child by i. Another example of such a family is formed by grids, where two different labels are used to distinguish edges to horizontal and vertical neighbours. (In our model there is no need to introduce special node labels that mark the boundaries of the grid, cf. [4, 2], as the labels of the edges incident to a node are available to our automata).

Automata. A k-head *graph-walking automaton with nested pebbles* is like its relative for trees in that it visits the nodes of a graph, each of its heads walking along the edges from node to node. It may check the labels of its current nodes, check whether such a node has an incident incoming/outgoing edge with a specific label (generalizing the concept of child number and the rankedness of Σ, respectively), check for the presence of a specific pebble (from a finite set), and it may then move each head along any edge in either direction choosing the edge based on its label. Moreover, it may drop and retrieve pebbles in a nested fashion, as before.

Formally we change the operations and tests available to the automaton. The moves up_i and $\mathrm{down}_{i,j}$ are replaced by the more symmetric pair $\mathrm{inmove}_{i,\sigma}$, $\mathrm{outmove}_{i,\sigma}$ to specify a move by the i-th head along an incoming (outgoing) edge with label σ. Likewise we replace the child-number test $\mathrm{chno}_{i,j}$ by the tests $\mathrm{inedge}_{i,\sigma}$ and $\mathrm{outedge}_{i,\sigma}$ on the existence of an incoming (outgoing) edge with label σ for the node currently under the i-th head.

Generally graphs do not have a distinguished node (like the root is a distinguished node for trees) and we change the definition of acceptance accordingly. A graph-walking automaton with nested pebbles accepts a given input graph if the automaton has an accepting computation (starting in the initial state, halting in an accepting state) when started with all its heads on *any* node of the input graph; initially and finally there should be no pebbles on the graph. We require that the existence or nonexistence of an accepting computation is independent of the chosen initial node, as this seems a natural condition, especially in the context of determinism. Note that not all automata satisfy this requirement.[7]

[7]Formally, we define $L(\mathcal{A})$ to consist of all graphs g over Σ such that $\mathrm{acc}(u)$ for every node u of g, where $\mathrm{acc}(u)$ means that $[q_0, u^k, \varepsilon] \vdash^*_{\mathcal{A},g} [q, \overline{v}, \varepsilon]$ for some halting configuration $[q, \overline{v}, \varepsilon]$ with $q \in A$ and $\overline{v}$ is a k-tuple

For trees over a ranked alphabet this definition is obviously equivalent to the old one as a tree-walking automaton can always move all its heads to the root. Disregarding pebbles, for grids the definition is equivalent to the k-head automaton of [2] (and for $k = 1$ to the 2-dimensional automaton of [4]).

In order to avoid an abundance of new notation, we keep the notation $\mathsf{DPW}^k\mathsf{A}$, etc., for families accepted by (deterministic) k-head automata (with nested pebbles) even in the more general context of graphs.

Logic. The first-order logic for graphs over the label alphabet Σ has atomic formulas $\mathrm{lab}_\sigma(x)$, $\sigma \in \Sigma$, for a node x with label σ, $\mathrm{edg}_\sigma(x, y)$, $\sigma \in \Sigma$, for an edge from x to y with label σ, and $x = y$.

As for automata, we keep the notation for families defined by our first-order logic with several variants of transitive closure. Note that we do not allow the predicate $x \leq y$, which makes the logic more like the local variant of first-order logic. As we have seen in Section 2, for trees this predicate is definable in first-order logic with (positive deterministic) transitive closure, and the families $\mathsf{FO{+}DTC}^k$, etc., of tree languages definable in first-order logic with transitive closure, do not change by this restriction.

For arbitrary families of graphs the computation of a graph-walking automaton with nested pebbles can be specified in first-order logic with transitive closure, like in Section 5.

Lemma 7.1. *For every family of graphs, and $k \geq 1$, $\mathsf{DPW}^k\mathsf{A} \subseteq \mathsf{FO{+}DTC}^k$.*

Proof. In the proof of Lemma 5.2, where we consider trees instead of graphs, no reference is made to properties particular to that domain. The only place where the proof given has to be adapted is where the accepting condition is phrased in the logic. Here any node must have an accepting computation when we start the automaton with all heads on that node, so we write $(\forall x)(\exists \overline{y}) \bigvee_{q \in A} \phi^{(n)\#}_{p,q}(x^k, \overline{y})$ where p is the initial state, the disjunction is taken over the set A of accepting states, and x^k is the k-tuple consisting of k copies of x. Note that we might as well take $\exists x$ instead of $\forall x$. □

The reverse inclusion needs an additional notion, the ability to search each graph in a given family of graphs.

Searchable Graphs. A family of graphs is *searchable* if there exists a (fixed) single-head deterministic graph-walking automaton with nested pebbles that, for each graph in the family, and each node of the graph, when started in that node in the initial state the automaton halts after completing a walk along the graph during which each node is visited at least once. Pebbles may be used in a nested fashion during the walk, as before. Thus the automaton serves as a *guide* for the family of graphs, and makes it possible to establish and traverse a total order of the nodes of the graph, generalizing the concept of preorder traversal used for trees. (In fact, the guide visits the same node perhaps several times, like in a preorder traversal on trees, but we will see that the unique first visit to a node can be recognized.) Note that the total order may depend on the node at which the guide starts its walk. Note also that if a family of graphs is searchable, then any subset of the family is searchable (by the same guide).

---

of nodes of g (and u^k is the k-tuple of k copies of u). Alternatively, we could require this for *some* node u of g, because, as discussed above, $\mathcal{A}$ is restricted to satisfy the requirement that for all graphs g over Σ and all nodes u_1, u_2 of g, if $\mathrm{acc}(u_1)$ then $\mathrm{acc}(u_2)$. We note, however, that this restriction is not essential for our results: it is easily shown that for searchable graphs (see below) the restriction can be dropped.

We give several elementary examples of this notion. Trees over a ranked alphabet form a searchable family: from any node we may walk to the root to start a preorder traversal through the tree. Similarly binary trees –which are basically ranked trees where every node has rank two, but where either child of a node can be missing– are a searchable graph family.

Unranked (ordered) trees, where there is no bound on the number of children of a node, are important in the theory of data representation using XML [47, 37, 57]. Here they are considered in their natural encoding as binary trees, where each node carries possibly two pointers (edges), one to its first child, one to its right sibling. In this way they are a searchable graph family. The (single-head) automata we obtain using this representation may move to the first child or to the next sibling of a node (and back), exactly as customary in the literature [47, 50] (albeit without pebbles). For an overview on logics for unranked trees, see [39].

Rectangular (directed) grids, edges pointing to the right or downwards, with edge labels distinguishing these two types of edges, form another example. This can be generalized to higher dimensional grids[8] [2].

All the above are superseded by acyclic, connected graphs with a single source (and the local restriction on edge labels). We can effectively turn those graphs into trees by placing an ordering on the labels and ignoring all incoming edges of a node except the edge which has minimal label among those edges. Graph-walking automata (with one head and no pebbles) on such graphs were considered in [36].

Obviously, a family of graphs such that, for some $\Sigma' \subseteq \Sigma$, the graphs induced by the Σ'-labelled edges belong to one of the above families, is searchable too. This includes the family of all 'ordered' graphs.

The above families are traversed in a rather standard way. In the next example we use a pebble to find and 'break' a cycle.

Example 7.2. Consider an ordinary binary tree, and take one of its leaves. If we identify this leaf with the root (i.e., we remove it and redirect the incoming edge to the root) we obtain a single directed cycle from which trees radiate, like charms hanging from a bracelet.

We argue that the family of these graphs is searchable. If started on a node in the cycle a graph-walking automaton may drop a pebble at that node. This effectively reduces the graph to a tree with the marked node as root, and a preorder traversal can be made.

We claim that a node on the cycle can be found, with the help of a single pebble, starting on an arbitrary node of the graph. Put the pebble on the initial node. Search the graph 'below' the pebble as if it is a tree; if indeed it is a tree, then we finally return to the pebble from below. Otherwise we are on the cycle, and we enter the leaf that has been identified with the root, and will find the pebble from above. Thus, in order to find a node on the cycle we repeat the above search, each time moving the pebble one node up while not on the cycle. □

Cyclic grids, or toruses, where the last node of each row has an edge to the first node of that row, and similarly for columns, can be searched using two pebbles. We search the grid row-by-row: the first pebble marks the position we start with (in order to stop when all rows are visited; we do not move this pebble during the traversal), the second pebble moves down in the first column to mark the position in which we started the row (in order to stop

[8] A d-dimensional grid has nodes $(x_1, \ldots, x_d)$ with $x_i \in \mathbb{N}$, $1 \leq x_i \leq k_i$, for certain $k_1, \ldots, k_d$. For each i it has i-labelled edges from $(x_1, \ldots, x_i, \ldots, x_d)$ to $(x_1, \ldots, x_i + 1, \ldots, x_d)$, provided $x_i < k_i$.

when we finish the row; we then move the pebble down to the next row until we meet the first pebble). This can be generalized to d pebbles in the d-dimensional case.[9]

It is an open question whether mazes (modelled as connected subgraphs of grids) form a searchable family of graphs, that is, with a single head and the help of nested pebbles. They cannot be searched without the help of pebbles, or with only a single (classic) pebble [12, 31]. According to [5] these graphs can be searched by single-head graph-walking automata with two (classic) pebbles which are unfortunately not used in a nested fashion, or using two heads (without pebbles).

The family of all graphs (over a given alphabet with at least two elements) is not searchable, not even with nonnested pebbles or with several heads. This follows from results of Cook and Rackoff [14]. The deterministic k-head graph-walking automaton with n pebbles is a special case of the jumping automaton of [14] with $k+n$ pebbles. A jumping automaton may move its pebbles along the edges of the graph, or move one of its pebbles to the location of another pebble ('jumping'). Retrievability of pebbles at a distance is a particular case of this jumping facility. Basically, [14, Theorem 4.9] states that the number of pebbles needed to visit all nodes of a d-dimensional torus grows as a function of d. Then [14, Theorem 4.13] concludes there is no jumping automaton that searches all graphs over a three letter alphabet, by coding arbitrary graph alphabets (like the d dimension edge labels for toruses) into a three letter alphabet. A slight adaptation of the proof (because of the graph model used) makes it valid in our setting, with two letters instead of three. For classic pebbles (and a single head) it is shown in [55] that the family of all planar graphs is not searchable.

For an overview of approaches to the exploration of mazes and graphs by finite automata see [24, 3].

For searchable graphs one can obtain the converse inclusion from our previous result, and translate logic into automata.

Theorem 7.3. *For every searchable family of graphs, and $k \geq 1$, $\mathsf{DPW}^k\mathsf{A} = \mathsf{FO{+}DTC}^k$.*

Proof. By the previous lemma it suffices to show $\mathsf{FO{+}DTC}^k \subseteq \mathsf{DPW}^k\mathsf{A}$ for searchable graphs, i.e., to reconsider the proof of Lemma 4.1. There is no need to return to a 'root' (we do not have one) to test acceptance by some automaton (the induction hypothesis used as a 'subroutine'); in the definition of acceptance we have required acceptance for any initial node, so we merely move all heads to the same position: drop a pebble, let each head search for it, and lift the pebble.[10]

At several points in that proof the automaton is required to make a systematic traversal through the input tree, or more specifically, to find the preorder successor of a given node, a function easily implemented for trees. At some other points we are supposed to move a head to a specific position marked by a pebble: the pebble again is found by a systematic traversal of the tree.

Here we instead use the 'guide', the automaton that makes the family searchable, as a subroutine to determine the successor of a node (and to make traversals of the graph). We have to mind some details, however. First, the guide may use its own pebbles during its computation, as is required to traverse a torus; leaving these pebbles on the graph would

[9]A d-dimensional torus is a d-dimensional grid with additional i-labelled edges from $(x_1, \ldots, k_i, \ldots, x_d)$ to $(x_1, \ldots, 1, \ldots, x_d)$.

[10]Dropping a temporary pebble is also a technique to show we may assume heads to be sensing, i.e., the automaton can check whether two heads occupy the same position.

block the removal of the other pebbles on the graph. Second, during the traversal of the graph the guide may visit the same node several times. In order to define the notion of successor properly, we have to distinguish a particular visit; for this we take the first visit to the node (as in the preorder of the nodes of a tree). Third, the order of visiting the nodes may depend on the node at which the guide is started, which makes the notion of successor undefined (think torus again).

Thus, we need to explain how the unique successor can be found for each node in the graph (except the last). First, at the very beginning of its computation, the automaton we construct for a given formula drops a pebble on the node where it is started (i.e., the node where all its heads are initially positioned; this may be any node of the graph). This pebble, let us call it 'the origin', remains at its position during the full computation to serve as a fixed position in the graph. Every successor will be determined in the order when starting the guide in the origin. Now the task of finding the successor of a (marked) node using a single head is executed as follows. First we use the guide (from wherever we are) to move the head to the origin. When the origin is reached, we 'reset' the guide (removing its pebbles from the graph, returning to its initial state) and again start a traversal of the graph looking for our marked node. When we reach the marked node, we continue running the guide until we reach a node that is visited for the first time: our (well-defined) successor. In order to check whether a node is visited for the first time we leave all pebbles of the guide at their position and drop a new pebble on the node. We then start a copy of the guide at the origin and run it until we find the new pebble. If the copy of the guide is in the same state and has the same positions for its pebbles as the original guide, then it visits the node for the first time: due to its determinism the guide cannot visit a node twice in the same configuration. After this test we stop the copy of the guide and retrieve its pebbles. We either have found the successor node, retrieve the pebbles of the guide and proceed as needed, or we lift the new pebble and continue the guide. If the guide halts without finding a first visit, the marked node was the last node in the order. □

As in Corollary 5.6, we have in the nondeterministic case $\mathsf{NPW}^k\mathsf{A} = \mathsf{FO{+}posTC}^k$ for every searchable family of graphs, and $k \geq 1$.

Multi-Searchable Graphs. It is open whether we can search a maze (a connected subgraph of a grid) with a single head using nested pebbles. However with two heads we can search a maze [5]. To cover this family we need to extend the notion of searchability: as the automata we are dealing with have k heads, it seems only fair to extend the guide with this commodity. A family of graphs is *k-searchable* if there is a deterministic guide as before, which now may have k heads. The guide starts its computation with all its heads on an arbitrary node of the graph, and performs a traversal of the graph. Recall that the formal model assumes that at most one head moves in each computation step of the guide; this ensures that at most one node is visited for the first time, thus uniquely defining the order of the nodes.

With this notion the results of the previous section can be extended to the larger class of k-searchable graphs. But there is a catch: we have to extend our automaton model with a *new instruction* $\langle p, \mathrm{jump}_i(x), q\rangle$ that moves a given head i to the position of a given pebble x (like the 'jumping' instruction from [14]). In the case of searchable graphs a single head can be moved to any pebble, just by running the guide using that head, searching for the pebble. For $k > 1$ however, in order to move a single head to a pebble, we need several heads, either introducing auxiliary heads (loosing the connection between number of heads

and arity of transitive closure) or loosing the position originally held by the heads needed for the search. (Even ignoring the problem of how to move all heads to the same initial position as needed for the guide.)

The suggested additional jump-instruction is rather natural if we think of pebbles as pointers as we did before. If the graph-walking automaton stores a finite number of 'addresses' of nodes, it can use a simple assignment to move a head to such a position. Adding this instruction is 'backward compatible': it will not change any of the previous results. On families of searchable graphs (which includes trees) the new instruction can be implemented as explained above.

Note that if a family is k'-searchable, then it is k-searchable for $k \geq k'$. Thus, the next result generalizes Theorem 7.3 (which for trees is Theorem 5.3).

Theorem 7.4. *Let the automaton model be extended with the jump-instructions, as discussed above. For every k-searchable family of graphs, $k \geq 1$, $\mathsf{DPW}^k\mathsf{A} = \mathsf{FO{+}DTC}^k$.*

Proof. For the inclusion $\mathsf{DPW}^k\mathsf{A} \subseteq \mathsf{FO{+}DTC}^k$ we note that the new instruction (move head i to pebble x) is directly expressible in the logic as we assume variables for the positions of heads and pebbles, cf. the proof of Lemma 5.2.

For the converse inclusion $\mathsf{FO{+}DTC}^k \subseteq \mathsf{DPW}^k\mathsf{A}$ we once more carefully inspect the proof of Lemma 4.1, or rather the proof of Theorem 7.3.

As observed in the latter proof, the guide is used for two purposes: (1) to find the successor of a node, and (2) to move a specific head to a specific pebble (and note that (2) is also used in (1): a kind of bootstrapping). For (2) we now use the new instruction, whereas for (1) we use the k-head guide. Note that to initiate the guide, all heads can be moved to the origin by the new instruction. Note also that before initiating the copy of the guide, pebbles should be dropped on all current head positions of the guide (to be able to compare them with those of the copy of the guide, and to restore them in case the successor has not been found). □

Since mazes are 2-searchable, we obtain that for mazes and $k \geq 2$, $\mathsf{DPW}^k\mathsf{A} = \mathsf{FO{+}DTC}^k$.

Again we obtain a similar result for the nondeterministic case: $\mathsf{NPW}^k\mathsf{A} = \mathsf{FO{+}posDTC}^k$ for every k-searchable family of graphs.

As discussed before Theorem 7.3, our k-head graph-walking automaton with nested pebbles is a special case of the jumping automaton of [14]; this includes the 'jumping' instruction of moving a head to the position of a pebble. Thus, by the results of [14], the family of all graphs (over $\Sigma \supseteq \{0,1\}$) is not k-searchable for any k.

According to Theorem 7.4, the existence of a k-head guide for a family of graphs entails that every language in $\mathsf{FO{+}DTC}^k$ can be implemented by a k-head graph-walking automaton with nested pebbles. This is obvious for the language L_0 defined by the formula $(\forall x)\,\mathrm{lab}_0(x)$: it can be checked directly by the guide whether every node visited has label 0. It is easy to see that this also works the other way around: assuming $\Sigma \supseteq \{0,1\}$, an automaton for L_0 must visit all nodes of any input graph that has only node label 0, since otherwise it also accepts that graph with the labels of unvisited nodes changed into 1. The automaton can be turned into a guide for the family by making it behave as if every node label equals 0. Here we require that the family is *node-label-insensitive*: membership of a graph in the family does not depend on the label of any node of the graph. For instance, the families of unranked trees, grids, and toruses we considered are node-label-insensitive. Thus, a node-label-insensitive family of graphs is k-searchable iff $L_0 \in \mathsf{DPW}^k\mathsf{A}$. Hence for such a family Theorem 7.4 also holds in the other direction: if $\mathcal{F}$ is a node-label-insensitive

family of graphs, then for every k, $\mathcal{F}$ is k-searchable iff $\mathsf{DPW}^k\mathsf{A} = \mathsf{FO{+}DTC}^k$. Moreover, $\mathcal{F}$ is k-searchable for some k iff $\mathsf{DPWA} = \mathsf{FO{+}DTC}$, where $\mathsf{DPWA} = \bigcup_{k\in\mathbb{N}} \mathsf{DPW}^k\mathsf{A}$, and $\mathsf{FO{+}DTC} = \bigcup_{k\in\mathbb{N}} \mathsf{FO{+}DTC}^k$. Moreover, this holds iff $\bigcup_{k\in\mathbb{N}} \mathsf{DW}^k\mathsf{A} = \mathsf{FO{+}DTC}$, assuming that the k-head graph-walking automaton (without pebbles) has an additional instruction to move one head to another (i.e., it is the jumping automaton of [14]).

Applying this to the family of all graphs, which is clearly node-label-insensitive, we note that the inclusions $\mathsf{DPW}^k\mathsf{A} \subset \mathsf{FO{+}DTC}^k$ (cf. Lemma 7.1) and $\mathsf{DPWA} \subset \mathsf{FO{+}DTC}$ are strict, the language L_0 being in the difference.

Pebbles left unturned. An obvious question that remains open is whether there exists a strict hierarchy in $\mathsf{DSPACE}(\log n)$ with respect to the number of heads used (allowing nested pebbles), or equivalently with respect to the arity of transitive closure. The same question can be asked for trees and grids. For arbitrary graphs, strictness of the hierarchies $\mathsf{FO{+}TC}^k$ and $\mathsf{FO{+}DTC}^k$ is shown in [29]. Some other unresolved questions were stated in Section 6, regarding single-head automata on trees. The question concerning the nature of our nested pebbles, whether the pointer model is more powerful than the classic model, remains interesting for multi-head automata. Another question is whether our results can be generalized to alternating automata and the alternating transitive closure operator of [33].

Acknowledgements

We thank Frank Neven and Thomas Schwentick, the authors of [48] (a first version appeared at ICALP 2000) where our result $\mathsf{FO{+}DTC}^1 = \mathsf{DPW}^1\mathsf{A}$ is cited, for their patience. We also thank them, Luc Segoufin, and David Ilcinkas for many useful comments. Finally, we thank the referees for their constructive remarks.

References

[1] A.V. Aho, J.D. Ullman. Translations on a context free grammar, Information and Control 19, 439–475, 1971. doi:10.1016/S0019-9958(71)90706-6

[2] Y. Bargury, J.A. Makowsky. The expressive power of transitive closure and 2-way multihead automata, in: Computer Science Logic, Proceedings 5th Workshop, CSL '91 (E. Boŕger, et al., eds.), Lecture Notes in Computer Science 626, 1–14, 1992. doi:10.1007/BFb0023754

[3] P. Beame, A. Borodin, P. Raghavan, W.L. Ruzzo, M. Tompa. A time-space tradeoff for undirected graph traversal by walking automata, SIAM Journal on Computing 28, 1051–1072, 1999. doi:10.1137/S0097539793282947

[4] M. Blum, C. Hewitt. Automata on a 2-dimensional tape, in: Proceedings 8th IEEE Symposium on Switching and Automata Theory, 155–160, 1967.

[5] M. Blum, D. Kozen. On the power of the compass (or, why mazes are easier to search than graphs), in: Proceedings 19th FOCS (Annual Symposium on Foundations of Computer Science), 132–142, 1978.

[6] M. Bojańczyk. 1-Bounded TWA cannot be determinized, in: Proceedings 23rd Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS'03 (P.K. Pandya, J. Radhakrishnan, eds.), Lecture Notes in Computer Science 2914, 62–73, 2003. doi:10.1007/b94618

[7] M. Bojańczyk, T. Colcombet. Tree-walking automata cannot be determinized, Theoretical Computer Science 350, 164–173, 2006. doi:10.1016/j.tcs.2005.10.031

[8] M. Bojańczyk, T. Colcombet. Tree-walking automata do not recognize all regular languages, in: Proceedings 37th ACM Symposium on Theory of Computing, STOC'05 (H.N. Gabow, R. Fagin, eds.), 234–243, 2006. doi:10.1145/1060590.1060626

[9] M. Bojańczyk, M. Samuelides, T. Schwentick, L. Segoufin. On the expressive power of pebble automata, in: Automata, Languages and Programming, 33rd International Colloquium, Proceedings Part I, ICALP '06 (M. Bugliesi, B. Preneel, V. Sassone and I. Wegener, eds.), Lecture Notes in Computer Science 4051, 157-168, 2006. doi:10.1007/11786986_15
[10] A. Brüggemann-Klein, D. Wood. Caterpillars: A context specification technique, Markup Languages 2, 81–106, 2000.
[11] J.R. Büchi. Weak second-order arithmetic and finite automata, Zeitschrift für mathematische Logik und Grundlagen der Mathematik 6, 66–92, 1960. doi:10.1002/malq.19600060105
[12] L. Budach. Automata and labyrinths, Mathematische Nachrichten 86, 195–282, 1978.
[13] B. ten Cate. The expressivity of XPath with transitive closure, in: Proceedings of the 25th ACM Symposium on Principles of Database Systems, PODS 2006, ACM Press, 328–337, 2006. doi:10.1145/1142351.1142398
[14] S.A. Cook, C.W. Rackoff. Space lower bounds for maze threadability on restricted machines, SIAM Journal on Computing 9, 636–652, 1980. doi:10.1137/0209048
[15] J. Doner. Tree acceptors and some of their applications, Journal of Computer and System Sciences 4, 406–451, 1970.
[16] H.-D. Ebbinghaus, J. Flum. *Finite Model Theory*, second edition, Perspectives in Mathematical Logic, Springer-Verlag, Berlin, 1999.
[17] C.C. Elgot. Decision problems of finite automata design and related arithmetics, Transactions of the American Mathematical Society 98, 21–52, 1961.
[18] J. Engelfriet. Context-free grammars with storage, Leiden University, Technical Report 86–11, 1986.
[19] J. Engelfriet, H.J. Hoogeboom. Tree-walking pebble automata, in: *Jewels are forever, contributions to Theoretical Computer Science in honor of Arto Salomaa* (J. Karhumäki et al., eds.), Springer-Verlag, 72–83, 1999.
[20] J. Engelfriet, H.J. Hoogeboom. Nested pebbles and transitive closure, in: Proceedings 23rd Annual Symposium on Theoretical Aspects of Computer Science, STACS 2006 (B. Durand, W. Thomas, eds.), Lecture Notes in Computer Science 3884, 477-488, 2006. doi:10.1007/11672142_39
[21] J. Engelfriet, H.J. Hoogeboom, J.-P. van Best. Trips on trees, Acta Cybernetica 14, 51–64, 1999.
[22] J. Engelfriet, S. Maneth. A comparison of pebble tree transducers with macro tree transducers, Acta Informatica 39, 613–698, 2003. doi:10.1007/s00236-003-0120-0
[23] J. Engelfriet, G. Rozenberg, G. Slutzki. Tree transducers, L systems, and two-way machines, Journal of Computer and System Sciences 20, 150–202, 1980. doi:10.1016/0022-0000(80)90058-6
[24] P. Fraigniaud, D. Ilcinkas, G. Peer, A. Pelc, D. Peleg. Graph exploration by a finite automaton, Theoretical Computer Science 345, 331–344, 2005. doi:10.1016/j.tcs.2005.07.014
[25] D. Giammarresi, A. Restivo. Two-dimensional languages, in: *Handbook of Formal Languages*, *Volume 3: Beyond Words* (G. Rozenberg, A. Salomaa, eds.), Chapter 4, Springer-Verlag, 1997.
[26] N. Globerman, D. Harel. Complexity results for two-way and multi-pebble automata and their logics, Theoretical Computer Science 169, 161–184, 1996. doi:10.1016/S0304-3975(96)00119-3
[27] E. Goris, M. Marx. Looping caterpillars, in: Proceedings 20th IEEE Symposium on Logic in Computer Science, LICS 2005, IEEE, 51–60, 2005. doi:10.1109/LICS.2005.24
[28] E. Grädel, G. McColm. On the power of deterministic transitive closure, Information and Computation 119, 129–135, 1995. doi:10.1006/inco.1995.1081
[29] M. Grohe. Arity hierarchies, Annals of Pure and Applied Logic 82, 103–163, 1996. doi:10.1016/0168-0072(95)00072-0
[30] J. Hartmanis. On non-determinancy in simple computing devices, Acta Informatica 1, 336–344, 1972. doi:10.1007/BF00289513
[31] F. Hoffmann. One pebble does not suffice to search plane labyrinths, in: Fundamentals of Computation Theory, FCT'81 (F. Gécseg, ed.), Lecture Notes in Computer Science 117, 433-444, 1981.
[32] O.H. Ibarra. Characterization of some tape and time complexity classes of Turing machines in terms of multihead and auxiliary stack automata, Journal of Computer and System Sciences 5, 88–117, 1971.
[33] N. Immerman. Languages that capture complexity classes, SIAM Journal on Computing 16, 760–778, 1987. doi:10.1137/0216051
[34] N. Immerman. Nondeterministic space is closed under complementation, SIAM Journal on Computing 17, 935–938, 1988. doi:10.1137/0217058

[35] N. Immerman. *Descriptive Complexity*, Graduate Texts in Computer Science, Springer-Verlag, New York, 1999.
[36] T. Kamimura, G. Slutzki. Parallel and two-way automata on directed ordered acyclic graphs, Information and Control 49, 10–51, 1981. doi:10.1016/S0019-9958(81)90438-1
[37] N. Klarlund, T. Schwentick, D. Suciu. XML: Model, Schemas, Types, Logics, and Queries, in: *Logics for Emerging Applications of Databases* (J. Chomicki, R. van der Meyden, G. Saake, eds.), Springer-Verlag, 1–41, 2004.
[38] S. C. Kleene. Representation of events in nerve nets and finite automata, in: *Automata Studies* (C.E. Shannon, J. McCarthy, eds.), Princeton University Press, 3–41, 1956.
[39] L. Libkin. Logics for unranked trees: an overview, Logical Methods in Computer Science 2, Issue 3, Paper 2, 2006. doi:10.2168/LMCS-2(3:2)2006
[40] O. Matz. Dot-depth, monadic quantifier alternation, and first-order closure over grids and pictures, Theoretical Computer Science 270, 1–70, 2002. doi:10.1016/S0304-3975(01)00277-8
[41] O. Matz, N. Schweikardt, W. Thomas. The monadic quantifier alternation hierarchy over grids and graphs, Information and Computation 179, 356–383, 2002. doi:10.1006/inco.2002.2955
[42] R. McNaughton, H. Yamada. Regular expressions and state graphs for automata, IRE Transactions on Electronic Computers 9, 39–47, 1960.
[43] T. Milo, D. Suciu, V. Vianu. Typechecking for XML transformers, Journal of Computer and System Sciences 66, 66–97, 2003. doi:10.1016/S0022-0000(02)00030-2
[44] B. Monien. Two-way multihead automata over a one-letter alphabet, RAIRO – Informatique Théorique et Applications 14, 67–82, 1980.
[45] A. Muscholl, M. Samuelides, L. Segoufin. Complementing deterministic tree-walking automata, Information Processing Letters 99, 33–39, 2006. doi:10.1016/j.ipl.2005.09.017
[46] L. Muzamel. Pebble alternating tree-walking automata and their recognizing power, Manuscript, 2006.
[47] F. Neven. Automata, logic, and XML, in: Computer Science Logic, 16th International Workshop, CSL 2002 (J.C. Bradfield, ed.), Lecture Notes in Computer Science 2471, 2–26, 2002.
[48] F. Neven, T. Schwentick. On the power of tree-walking automata, Information and Computation 183, 86–103, 2003. doi:10.1016/S0890-5401(03)00013-0
[49] F. Neven, T. Schwentick, V. Vianu. Finite state machines for strings over infinite alphabets, ACM Transactions on Computational Logic 5, 403–435, 2004. doi:10.1145/1013560.1013562
[50] A. Okhotin, K. Salomaa, M. Domaratzki. One-visit caterpillar tree automata, Fundamenta Informaticae 52, 2002, 361–375.
[51] H. Petersen. The equivalence of pebbles and sensing heads for finite automata, in: Proceedings 11th Symposium Fundamentals of Computation Theory, FCT'97, Lecture Notes in Computer Science 1279, 400–410, 1997. doi:10.1007/BFb0036201
[52] A. Potthoff. Logische Klassifizierung regulärer Baumsprachen, PhD thesis, Institut für Informatik und Praktische Mathematik, Bericht Nr. 9410, Universität Kiel, 1994.
[53] M.O. Rabin, D. Scott. Finite automata and their decision problems, IBM Journal of Research and Development 3, 114–125, 1959.
Also in: *Sequential Machines: Selected Papers* (E.F. Moore, ed.), Addison-Wesley, Reading, MA, 1964, pp. 63–91.
http://www.research.ibm.com/journal/rd/032/ibmrd0302C.pdf
[54] R.W. Ritchie, F.N. Springsteel. Language recognition by marking automata, Information and Control 20, 313–330, 1972. doi:10.1016/S0019-9958(72)90205-7
[55] H.A. Rollik. Automaten in planaren Graphen, Acta Informatica 13, 287–298, 1980. doi:10.1007/BF00288647
[56] A. Rosenfeld. *Picture Languages*, Academic Press, New York, 1979.
[57] T. Schwentick. Automata for XML – a survey, Journal of Computer and System Sciences 73, 289–315, 2007. doi:10.1016/j.jcss.2006.10.003
[58] J.C. Shepherdson. The reduction of two-way automata to one-way automata, IBM Journal of Research and Development 3, 198–200, 1959.
Also in: *Sequential Machines: Selected Papers* (E.F. Moore, ed.), Addison-Wesley, Reading, MA, 1964, pp. 92–97.
http://www.research.ibm.com/journal/rd/032/ibmrd0302P.pdf

[59] M. Sipser. Halting space-bounded computations, Theoretical Computer Science 10, 335–338, 1980. doi:10.1016/0304-3975(80)90053-5

[60] G. Slutzki. Alternating tree automata, Theoretical Computer Science 41, 305–318, 1985. doi:10.1016/0304-3975(85)90077-5

[61] R. Szelepcsényi. The method of forced enumeration for nondeterministic automata, Acta Informatica 26, 279–284, 1988. doi:10.1007/BF00299636

[62] J.W. Thatcher, J.B. Wright. Generalized finite automata theory with an application to a decision problem of second-order logic, Mathematical Systems Theory 2, 57–81, 1968. doi:10.1007/BF01691346

[63] W. Thomas. Languages, automata, and logic, in: *Handbook of Formal Languages*, *Volume 3: Beyond Words* (G. Rozenberg, A. Salomaa, eds.), Chapter 7, Springer-Verlag, 1997.

[64] B.A. Trakhtenbrot. Finite automata and the logic of one-place predicates, Siberian Mathematical Journal 3, 103–131, 1962. (in Russian). English translation: American Mathematical Society Translations, Series 2 59, 23–55, 1966.

[65] V. Vianu. A Web Odyssey: from Codd to XML, in: Proceedings of the 20th ACM Symposium on Principles of Database Systems, PODS 2001, ACM Press, 1–15, 2001. doi:10.1145/375551.375554

[66] H. Vogler (ed.). International Workshop on Grammars, Automata, and Logic on Graphs and Trees, Technical Report TUD-FI99-01, Dresden University of Technology, March 1999.

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/2.0/ or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.